Fair Packet Forwarding in MANETs with Anonymous Stations: A Game-Theoretic Approach

Jerzy Konorski[1]

[1]Gdansk University of Technology
ul. Narutowicza 11/12, 80-952 Gdansk, Poland
jekon@eti.pg.gda.pl

Abstract: A station of a mobile ad-hoc network (MANET) may selfishly refuse to forward transit packets as it shortens the battery life and takes up a portion of the bandwidth that could be used for source packets. Due to a high degree of station anonymity, selfishness meets with little punishment. The well-known watchdog mechanism can be used to check if an adjacent station forwards packets. We point out that a watchdog may be unable to tell source from transit packets, which enables undetectable manipulation of local congestion controls in selfish stations. We allow each station to set its source packet admission threshold so as to maximise a throughput- and reputation-related payoff. The nature of possible Nash equilibria of the resulting noncooperative game are examined for a generic model of packet forwarding and symmetric traffic flows. A novel packet forwarding protocol called F^3T is proposed and the payoffs it yields are approximately analysed.

1 Introduction

A mobile ad-hoc network (MANET) consists of a number of mobile stations exchanging data packets over one or more wireless channels. MANETs rely on self-organisation rather than central administration. Owing to the falling prices of wireless network equipment as well as the advances in protocol design, MANETs have abandoned their traditional niche of military and emergency communications and increasingly enter the field of civilian data services [6]. Logically, a MANET can be visualised as a (time-varying) adjacency graph, station *m* being adjacent to station *n* if it remains within the latter's hearing range. Each station acts both as a mobile user terminal and a packet forwarder. As the former, it injects *source packets* into the network and absorbs *destination packets* therefrom; as the latter, it forwards *transit packets* on behalf of currently non-adjacent pairs of stations. Forwarding transit packets is a dual liability: it shortens the station's battery life and takes up a portion of the channel bandwidth it could use to transmit source packets. MANETs allow a high degree of station anonymity and so refusal to forward transit packets may meet with little punishment. One can therefore envisage various types of station misbehaviour; [11] presents a comprehensive taxonomy.

We focus on *selfish* behaviour whereby stations try to reap some undue benefits (as distinct from *cooperative* behaviour and from *malicious* behaviour meant just to do some damage unto others). For a taxonomy of selfish behaviour in MANETs, see [10]. Cooperative behaviour cannot be enforced in a MANET other than by some incentive-based mechanisms. Buttyan and Hubaux [2] propose a virtual currency called *nuglets* that a station earns by forwarding transit packets and then uses to buy a similar service from other stations. Marti *et al*. [9] propose to equip a station with a *watchdog* mechanism which listens to adjacent stations' transmissions and checks if they perform forwarding. A number of recent papers adopts game theory, whereeach player (station) sets her own strategy at will, but the received payoff also depends on the other players' strategies. The play often reaches a Nash equilibrium (NE) from which no player wants to deviate [5]. Michiardi and Molva [10] incorporate a measure of reputation into the payoffs so that rational players forward transit packets to avoid being excluded from existing routing paths. Urpi *et al.* [14] and Srinivasan *et al.* [12] relate the payoffs to throughput efficiency and battery consumption. Zhong *et al.* [15] show that honesty in handling virtual currency can be made a payoff maximising strategy. In the approach of Felegyhazi *et al*. [4], each station sets its own level of cooperation based on its current perception of other stations' levels.

We examine the watchdog approach and argue that some of its weaknesses listed in [9] are not fundamental. However, we point to an unlisted one: being unable to tell source from transit packets, a watchdog is also unable to decide whether an adjacent station is misbehaving or it is backlogged due to heavy transit traffic. Undetectable selfish behaviour then consists in over-admittance of source packets. We address this issue in a game-theoretic framework. Next we propose a packet forwarding protocol called F^3T under which the NE of the underlying game prescribes fair and throughput-efficient settings of local congestion controls.

In Sec. 2 we formulate the network model and explain the nature of undetectable selfish behaviour. In Sec. 3 we define a noncooperative congestion control game and discuss its outcomes. In Sec. 4 we describe and approximately analyse the F^3T protocol. In Sec. 5 we discuss the relevance of proper configuration of F^3T from a game-theoretic perspective. Sec. 6 concludes the paper.

2 Network and packet forwarding model

We assume that the stations use omnidirectional antennae so that the watchdogs can hear all adjacent stations' transmissions. The adjacency graph is assumed bidirectional. MAC and multihop routing protocols are not relevant to our considerations and will not be specified. MAC addresses need not be trustworthy. Data privacy and station anonymity can be achieved via a public-key cryptosystem, such as RSA (with off-line encryption and decryption), and a public hash function such as SHA-1 or MD-5 [13]. Public keys need not be permanent or unique per station. Packet forwarding can now be outlined as follows (Fig. 1):

- a pair of adjacent stations, n and m, establish a *neighbourhood relationship* by exchanging their public keys, key_n and key_m, and routing tables,
- to transfer a packet to a destination station d, a source station n first looks up the next-hop neighbour station m in the routing table, then uses key_d to encrypt the packet body along with key_n, next appends $h_n = hash(key_n)$, $h_m = hash(key_m)$ and $h_d = hash(key_d)$ and finally transmits the packet,
- if the packet is received error-free, station m acks it and compares h_n, h_m and h_d with locally stored hashes of public keys to check that it has a neighbourhood relationship with station n and to recognise itself as the receiver and (possibly) destination; if $m = d$ is detected, the packet body is decrypted using the private key key_m^{-1}, otherwise h_m is replaced by $h_l = hash(key_l)$ with l determined by station m's routing table, and the forwarding continues,
- if $m \neq d$, station n performs a *watchdog check*: upon reception of station m's transmission with h_l appended, it compares the packet body with a copy it has retained to check that the packet has indeed been forwarded by station m,
- based on the check statistics, a neighbourhood relationship may be terminated.

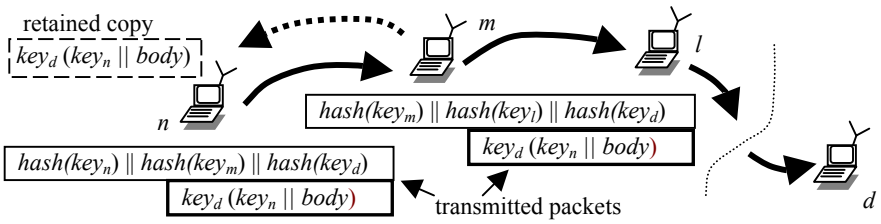

Fig. 1. Anonymous forwarding with public keys and a hash function

Packet collisions on a wireless channel may cause some ambiguity e.g., the watchdog at station n may not receive a packet being forwarded by station m or may be uncertain if the packet has been received at station l [9]. This danger is not serious if a powerful enough MAC protocol is employed e.g., CSMA/CA with RTS/CTS or multiple-channel CDMA. Also, most types of selfish behaviour can be countered under the above model. Refusal to forward a packet claiming transmission errors is counterproductive since the neighbour stations are likely to monitor the channel quality. Assumption of a new identity (new key_m) with a 'clean record' can be countered by requiring an initial silent period for each neighbourhood relationship. False deleting a neighbour from the routing table would not pay off if 'rich enough' routing tables were required to maintain a neighbourhood relationship. Finally, collusion between a pair of neighbours tolerating each other's misbehaviour is risky as either of them might use it for not forwarding the other's packets.

Note that the watchdog at station n must allow station m to transmit a number of packets prior to one being listened for, up to a public-knowledge deadline B beyond which a *failed check* occurs. Station m can legitimately refuse to receive a packet (via a Receive-Not-Ready frame) claiming a current backlog in excess of B. To prove its

claim, station m appends to the RNR frame the hashes of B backlogged packets, to be subsequently listened for at station n. Pretending to have a large backlog implies appending fake hashes and subsequently transmitting dummy packets which compute to the same hashes. This could only be productive if the dummy packets were unusually short; an obvious remedy is to define some minimum packet size.

The encryption of a source station's public key (for reasons of privacy and anonymity) precludes the watchdog at station n from distinguishing between station m's source and transit packets. This leaves a possibility of undetectable selfish behaviour. Namely, the necessity to keep the backlog low mandates a local congestion control mechanism at each station. E.g., a simple Drop-and-Throttle (D&T) mechanism [7] permits a station to admit source packets only if its current backlog is below a, where a is called the *D&T threshold*. Since a is set locally, nothing stops station m from unrestrained admission of source packets and subsequent issuing of RNR frames. Such misbehaviour will go unnoticed as the watchdog at station n, listening to station m's transmissions, is unable to tell source packets from transit packets; legitimate refusal to receive packets is indistinguishable from selfish behaviour. A packet forwarding protocol is therefore needed that offers incentives to set a so as to achieve fair and efficient use of the channel bandwidth.

3 Game-theoretic model

In this section we describe a noncooperative 'D&T game' that arises when each station sets its D&T threshold so as to maximise the local source packet admission rate while keeping failed checks at neighbour stations tolerably rare. Following game theory [5], we presume that a set of selfish stations reach a Nash equilibrium. We show that even if the traffic flows are symmetric and the D&T thresholds are initially identical, the outcome at equilibrium may be unfair to some stations.

3.1 The D&T game

Let us view all stations as players in a nonzero-sum game. Station n's feasible actions are the values of D&T threshold it sets locally. A *D&T threshold profile* has the form $[a_n\,(a)_{-n}]$, where a_n is station n's D&T threshold and $(a)_{-n} = (a_m, m \neq n)$ is the *opponent profile*. The payoff to any station n is determined by the D&T threshold profile. We define two payoff components:

- a throughput measure $S[a_n\,(a)_{-n}]$ – the local source packet admittance rate, and
- a deadline violation measure $V[a_n\,(a)_{-n}]$ – the station n-related rate of failed checks at a neighbour station i.e., the rate of reception of packets in whose presence station n will have transmitted at least B other packets.

The former component gives incentives to increase a_n, whereas the latter gives incentives to keep a_n moderate lest station n's selfish behaviour be detected. We take

$$V[a_n\,(a)_{-n}] \geq V^* \qquad (1)$$

as the condition of termination of all neighbourhood relationships involving station n, where V^* is a public-knowledge tolerance level. V^* should be set distinctly above the station malfunction rate. Also, it should be large enough for (1) to be detected with statistical credibility. V^* should be upper bounded in relation to the average route length; e.g., with 5-hop routes and $V^* \leq 10^{-3}$, over 99.5% of traffic reaches destination. To include reputation effects in the payoffs we assume that

(i) Any station is interested in maintaining all its neighbourhood relationships for which (1) is false and none for which it is true.

Station n therefore wants to maximise

$$payoff[a_n\,(a)_{-n}] = \begin{cases} S[a_n\,(a)_{-n}], \text{ if } V[a_n\,(a)_{-n}] < V^* \\ 0, \text{ otherwise} \end{cases} \tag{2}$$

A *Nash equilibrium* (NE) is a D&T threshold profile $[a_n^0(a^0)_{-n}]$ such that

$$payoff[a_n^0(a^0)_{-n}] \geq payoff[a_n(a^0)_{-n}] \text{ for all } n \text{ and any } a_n \tag{3}$$

i.e., no station has incentives to change its D&T threshold unilaterally.

The exact form of the payoffs (2) is determined by a packet forwarding protocol. The following assumption states the payoff structure needed for the D&T game to be nontrivial. The first two parts imply that when increasing its D&T threshold, station n faces a conflict between the increased $S[a_n\,(a)_{-n}]$ and $V[a_n\,(a)_{-n}]$; the third part implies that unrestrained admission of source packets does not pay off:

(ii) If D&T threshold profiles are ordered in the sense of vector inequality then
- $S[a_n\,(a)_{-n}]$ increases in a_n and decreases in $(a)_{-n}$,
- $V[a_n\,(a)_{-n}]$ increases in both a_n and $(a)_{-n}$, and
- $V[a_n\,(a)_{-n}] \geq V^*$ for a large enough $[a_n\,(a)_{-n}]$.

In response to changes in $(a)_{-n}$, station n can adjust its D&T threshold. To capture the game dynamics while keeping the model simple, we assume that

(iii) The underlying D&T threshold adjustment mechanism is
- *locally adaptive* i.e., a_n is adjusted based on observed changes in $payoff[a_n\,(a)_{-n}]$,
- *gradual* i.e., a change of a_n by $\pm\Delta$ causes a payoff change to any other station equivalent of Δ consecutive changes of a_n by ± 1, and
- *prompt* i.e., any other station becomes aware of and can react to each of these Δ changes before the next one takes effect.

The first part reflects the distributed nature of ad-hoc network protocols, whereas the second part is motivated by the fact that $S[a_n\,(a)_{-n}]$ and $V[a_n\,(a)_{-n}]$ are medium-term statistics, so do not change abruptly. The third part implies unit changes of the D&T thresholds: each station n moves sequentially; each move consists in changing a_n by

±1 and immediately yields payoffs corresponding to the new D&T threshold profile; and no station lags more than one move behind the other stations.

3.2 Nash equilibria of a symmetric D&T game

Suppose that the traffic flows are symmetric and a D&T threshold profile $[a_0\,(a_0 \ldots a_0)]$ currently prevails, yielding fair nonzero payoffs. The following proposition categorises possible Nash equilibria of the D&T game (see [8] for proof and Sec. 5 for illustration).

Under assumptions (ii) and (iii), a symmetric D&T threshold profile [a (a … a)] will eventually be reached whereupon

- *each station receives a nonzero payoff and has no incentive to change its D&T threshold (a symmetric efficient NE), or*
- *all neighbourhood relationships will be terminated since each station receives a zero payoff, but has no incentive to change its D&T threshold (a symmetric inefficient NE), or*
- *a timing game (a 'war of preemption' or 'war of attrition') starts, leading to an asymmetric NE with unfair payoffs.*

Remarks:

- In a timing game, a player moves at most once and initially all players have incentives to move. In a 'war of preemption,' moving early yields higher payoffs than moving late or not at all, whereas in a 'war of attrition,' moving late or not at all yields higher payoffs.
- Of the above outcomes, the first one (a symmetric efficient NE) is the only desirable; in the following sections we will show how it can be attained by proper design and configuration of the packet forwarding protocol.
- Fairness is somewhat difficult to define in the case of asymmetric traffic flows; for the purpose of this paper we consider a packet forwarding protocol satisfactory if it yields a symmetric efficient NE for the symmetric D&T game.

4 F^{3}T protocol

Any packet forwarding protocol ought to include mechanisms of 1) provably legitimate refusal to receive packets (otherwise a station may be unduly punished for increased transit traffic), and 2) provably legitimate override of 1) (otherwise $V[a_n\,(a)_{-n}]$ need not increase in a_n or $(a)_{-n}$, contrary to assumption (ii)). We present a protocol called *Fair Forwarding with Forced Transmissions* (F^{3}T) and approximately analyse the D&T game payoffs under F^{3}T for a symmetric network model.

4.1 Protocol description

Depending on its current backlog x, a station operates in the NORMAL mode (when $x < e$) and the CONGESTED mode (when $x \geq e$). The parameter e is public knowledge. In the CONGESTED mode, station n can legitimately refuse to receive packets. This it does by announcing the hashes of e backlogged packets, appended to an RNR frame or a transmitted packet. Let the current backlog at a neighbour station m be y. If $y < e$, station m suspends further packet transmissions to station n until $x < e$, as announced by the latter via a Receive-Ready frame or a suitable indication in a transmitted packet. On the other hand, if $y \geq e$ and station m has a packet ready for station n, it *forces* a packet transmission and appends to it the hashes of e backlogged packets as proof of the CONGESTED mode. Thus a CONGESTED station requests that inbound packet transmissions be suspended, which NORMAL neighbour stations comply with and CONGESTED ones disregard. Note that the protocol operation for $e > B$ is the same as for $e = 1$.

A backlogged packet at station n whose next-hop station is m remains enqueued if

- $x < e$ and $y \geq e$ i.e., a packet transmission to station m cannot be forced, or
- no channel is available for a packet transmission to station m, or
- some packet received prior to the packet in question remains enqueued.

Thus packets at a station form a common FIFO queue regardless of the selected next-hop stations. FIFO queuing is known to reduce channel utilisation; on the other hand it enables correct F^{3}T protocol operation in the following way:

- with per next-hop station queues it would be unclear whether a failed check is due to selfish behaviour or the fact that the neighbour station keeps transmitting packets from other queues,
- a packet received when $x = B$ is certain to cause a failed check; as such it can be immediately discarded instead of unproductively increasing the backlog, and
- a claim of the CONGESTED mode at a neighbour station can be verified within a definite time horizon by comparing the hashes of e subsequent packets transmitted by that station with the previously received hashes.

A neighbourhood relationship is terminated if a claim of the CONGESTED mode is not verified (the other condition is (1)). The goal is to configure B, V^* and e such that the D&T game has a symmetric efficient NE.

4.2 D&T game payoffs under F^{3}T

Assumption (iii) in Sec. 3.1 restricts our interest to payoffs to D&T threshold profiles of the form $[a'(a \ \ldots \ a)]$ with $|a' - a| \leq 1$. These will be calculated assuming that:

(A1) all stations synchronise to fixed-size time slots, a slot accommodating a packet transmission along with related acks, RNR and RR frames,
(A2) each station has M neighbour stations,
(A3) the average source-to-destination path length is H hops,
(A4) the next-hop station for a backlogged packet is selected at random,

(A5) the network operates under heavy load i.e., in each slot a station admits as many source packets as its D&T threshold permits, and
(A6) a station can simultaneously and error-free transmit to and receive from all its neighbour stations, at most one packet per slot per neighbour station.

Assumption (A4) simplifies the calculation; it reflects, in a somewhat exaggerated way, the path variability in ad-hoc networks. Assumption (A5) factors out traffic generation characteristics. Finally, assumption (A6) implies a powerful multipacket reception scheme e.g., CDMA. We stick to this assumption to avoid shifting the focus from packet forwarding to multiple access and physical transmission.

Following the 'isolated node' approach [1], we shall focus upon an arbitrarily chosen station n, where the current backlog x will be modelled as a homogeneous Markov chain. The transition probabilities depend on the current backlog $y_1, ..., y_M$ at the neighbour stations $n_1, ..., n_M$. The approximation consists in regarding $y_1, ..., y_M$ in each slot as drawn from the steady-state probability distribution ($p(x)$, $0 \le x \le B$) of the above Markov chain. This leads to a fixed-point relationship of the form $p(\cdot) = f[p(\cdot)]$, which can be solved iteratively for $p(\cdot)$. Let $\mathrm{X}^{(s)}$ and $\mathrm{Y}_m^{(s)}$ denote the backlog at station n and n_m, respectively, at the start of the s^{th} slot. The $\mathrm{Y}_m^{(s)}$ will be treated as iid with respect to s and m; let $p_{\mathrm{Y}}(y) = \Pr[\mathrm{Y}_m^{(s)} = y]$ for $0 \le y \le B$ and $1 \le m \le M$. We will express the transition probabilities for $\mathrm{X}^{(s)}$ through $p_{\mathrm{Y}}(\cdot)$.

Given $\mathrm{X}^{(s)} = x$ and $(\mathrm{Y}_1^{(s)} ... \mathrm{Y}_M^{(s)}) = (y_1 ... y_M) = \mathbf{y}$, denote by $\mathrm{T}|_{x,\mathbf{y}}$ and $\mathrm{R}_m|_{x,\mathbf{y}}$ respectively the random number (between 0 and $\min[x, M]$) of packet transmissions out of station n and the number (0 or 1) of non-destination packets received from station n_m in the s^{th} slot. Let $P_{\mathrm{Y}}(e) = \sum_{0 \le y < e} p_{\mathrm{Y}}(y)$. Recalling assumption (A6) and the conditions for a packet to remain enqueued (Sec. 4.1), one has

$$\Pr[\mathrm{T}|_{x,\mathbf{y}} \ge k] = \begin{cases} 0, \text{ if } k > \min[x, M] \\ P_{\mathrm{Y}}^k(e) \cdot \prod_{0 \le j \le k-1} (1 - \frac{j}{M}), \text{ if } k \le \min[x, M] \text{ and } x < e \\ \prod_{0 \le j \le k-1} (1 - \frac{j}{M}), \text{ if } k \le \min[x, M] \text{ and } x \ge e \end{cases} \quad (4)$$

$$\Pr[\mathrm{R}_m|_{x,\mathbf{y}} = 1] = \begin{cases} 0, \text{ if } y_m < e \text{ and } x \ge e \\ (1 - \frac{1}{H}) \cdot \frac{1}{M} \cdot \sum_{0 \le j < \min[y_m, M]} \prod_{0 \le i \le j} P_{\mathrm{Y}}^j(e), \text{ otherwise} \end{cases} \quad (5)$$

The second part accounts for various positions (1^{st},..., $\min[y_m, M]^{\text{th}}$) the packet to be transmitted to station n may occupy in station n_m's FIFO queue. The random variable

$$\mathrm{I}|_{x,\mathbf{y}} = \sum_{1 \le m \le M} \mathrm{R}_m|_{x,\mathbf{y}} - \mathrm{T}|_{x,\mathbf{y}} \quad (6)$$

represents the net influx of transit packets at station n per slot. Its probability distribution is obtainable via (4) and (5) since $\mathrm{T}|_{x,\mathbf{y}}$ and $\mathrm{R}_m|_{x,\mathbf{y}}$ are independent. Unconditioning on **y** gives

$$\Pr[\mathrm{I}|_x = i] = \sum_{0 \le y_1,\ldots,y_M \le B\}} p_{\mathrm{Y}}(y_1)\ldots p_{\mathrm{Y}}(y_M) \cdot \Pr[\mathrm{I}|_{x,\mathbf{y}} = i]\,. \tag{7}$$

Note that there are $(B+1)^M$ summands in (7), each of which involves numerical inversion of a probability generating function for (6); this makes (7) the most tedious part of the calculation. Given $\mathrm{X}^{(s)} = x$ one has

$$\mathrm{X}^{(s+1)} = \min[B,\, x + (a-x)^+ + \mathrm{I}|_x]\,, \tag{8}$$

where $(a–x)^+ = \max[0,\, a–x]$ is the number of source packets admitted in the s^{th} slot. Hence, the calculation of the transition probabilities of interest as well as the steady-state probabilities $p(x) = \lim_{s\to\infty}\Pr[\mathrm{X}^{(s)} = x]$ is straightforward, assuming that $p_{\mathrm{Y}}(\cdot)$ is known. Since by symmetry $p(\cdot) = p_{\mathrm{Y}}(\cdot)$, one can calculate $p(\cdot)$ based on an assumed $p_{\mathrm{Y}}(\cdot)$, and in the next iteration substitute $p(\cdot)$ for $p_{\mathrm{Y}}(\cdot)$ until the two differ insignificantly. Thereupon one obtains

$$S[a(a\ldots a)] = \sum_{0\le x\le a} p(x)\cdot(a-x)^+ \tag{9}$$

$$V[a(a\ldots a)] = \frac{1}{M}\cdot\sum_{0\le x\le B} p(x)\sum_{i>B}(i-B)^+ \cdot \Pr[x+(a-x)^+ + \mathrm{I}|_x = i] \tag{10}$$

To calculate *payoff*$[a'\,(a\ \ldots\ a)]$, put $a := a'$ in (9) and (10) while retaining $p(\cdot)$.

5 F^{3}T configuration and performance

The approximate analysis outlined in Sec. 4.2 assumes a rather idealised network model and only yields to numerical calculation. Yet its results are instructive as they capture the possible outcomes of the D&T game under F^{3}T. Provided that B, e and V^* are set properly, increasing a_n unilaterally backfires in terms of V: a larger backlog at station n initially reduces the local packet reception rate, but ultimately drives the neighbour stations CONGESTED and causes them to force packet transmissions into n. Without incentives to increase a, the adjustment mechanism mentioned in assumption (iii) of Sec. 3.1 remains dormant and a fair and throughput-efficient D&T threshold profile persists. The results also illustrate the difference in performance prediction compared to the classical cooperative paradigm.

A series of numerical experiments confirmed the validity of assumption (ii) of Sec. 3.1 under F^{3}T. For various a and $a' = a–1$, a and $a+1$, *payoff*$[a'\,(a\ \ldots\ a)]$ was calculated from (9) and (10) as a percentage of the maximum attainable value M/H. Sample results are depicted in Fig. 2a-c, assuming $B = 8$, $M = 3$, $H = 5$ and $V^* = 10^{-3}$.

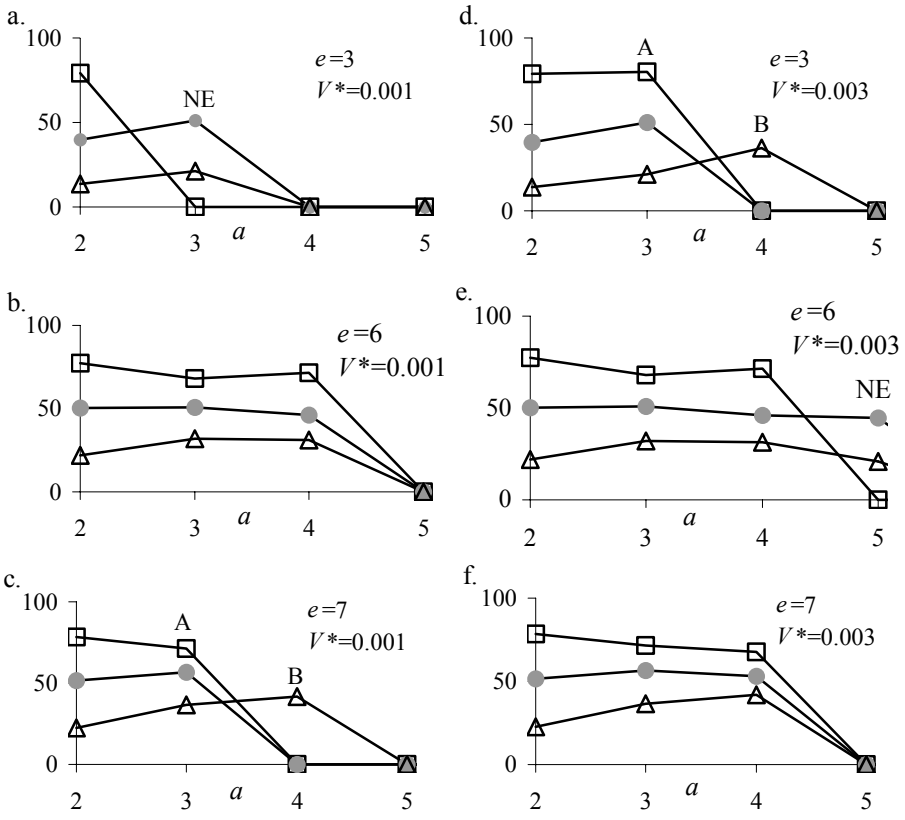

Fig. 2. *payoff*[a' (a…a)] (%) under F^{3}T (□: a'=a+1, **O**: a'=a, **Δ**: a'=a−1)

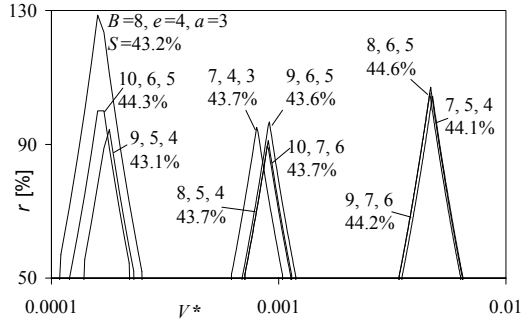

Fig. 3. Robustness of symmetric efficient Nash equilibria under F^{3}T

One sees that e =3 yields a symmetric efficient NE with S[3 (3 ... 3)] = 51.2%: no station has incentives to increase or decrease its D&T threshold. With e = 6, at the profile [4 (4 ... 4)] all stations have incentives to increase a, ending up at the profile [5 (5 ... 5)] with no incentive to move further. This leads to the termination of all neighbourhood relationships. With e = 7, at the profile [3 (3 ... 3)] all stations benefit from moving to a = 4. Supposing they do not do so at one time, a 'war of preemption' follows: early movers end up at point A, while late movers end up at B with no incentive to move further (since the payoff corresponding to [3 (4 ... 4)] exceeds that corresponding to [4 (4 ... 4)]). Had all the stations moved from a = 3 to a = 4 simultaneously, they would have to consider retreating to a = 3, but early movers now end up at B and late movers end up at A. This is in effect a 'war of attrition.' The (unfair) payoffs thus may range from 71.4% (at A) to 41.8% (at B). Note that fair payoffs of 56.7% could be attained for e = 7 if the stations were cooperative and stuck to a = 3.

The outcome of the D&T game varies with V*, cf. Fig. 2d-f. Given B and e, a choice of V* yielding a symmetric efficient NE is always possible, though may result in different *robustness*. Suppose that a NE occurs at [a (a ... a)]. Then it must be that V[a (a ... a)] < V* and V[a+1 (a ... a)] ≥ V* (cf. Fig. 2a,e). Since V is a statistical average and V* is typically low in magnitude, the values of V[a (a ... a)], V* and V[a+1 (a ... a)] ought to be quite distinct so that each station can avoid accidental departure from the NE. Therefore, the lesser of their relative differences, $r = \min\{V^*/V[a\ (a\ ...\ a)]-1,\ V[a+1\ (a\ ...\ a)]/V^*-1\}$, measures the robustness of the NE. Fig. 3 shows the ranges of V* yielding a symmetric efficient NE with $r \geq 50\%$ (the corresponding B, e, a and S are indicated). Thus that a proper setting of B and e guarantees a desirable and robust outcome for various magnitudes of V*. However, the received payoffs are relatively invariant and under 45%.

5 Conclusion and future research

We have defined and analysed a noncooperative 'D&T game' played by anonymous MANET stations. For symmetric flows and a novel protocol called F^{3}T we have shown how the stations can be given incentives to reach a fair and throughput-efficient NE in terms of local congestion control settings. Possible improvements and extensions of the presented work include:

- simulative analysis of asymmetric D&T games under various dynamic scenarios,
- detailed calculation of the transmission overhead related to F^{3}T operation and the employed hash function,
- other e.g., rate-based mechanisms of source packet admission,
- single-channel MAC protocols e.g., IEEE 802.11,
- per next-hop station queueing to improve the channel utilisation (a suitable extension of F^{3}T is possible under source routing schemes e.g., DSR [3]), and
- individual negotiation of B, e and V* for each neighbourhood relationship.

References

1. Agnew, G. B., Mark, J. W.: Performance Modeling for Communication Networks at a Switching Node. IEEE Trans. Comm. COM-32 (1984) 902-910
2. Buttyan, L., Hubaux, J. P.: Nuglets: A Virtual Currency to Stimulate Cooperation in Self-Organised Mobile Ad-Hoc Networks. Tech. Rep. DSC/2001/001, Swiss Federal Institute of Technology (2001)
3. Broch, J., Johnson, D., Maltz, D.: The Dynamic Source Routing Protocol for Mobile Ad Hoc Networks. IETF Internet Draft (1998)
4. Felegyhazi, M., Buttyan L., Hubaux, J. P.: Equilibrium Analysis of Packet Forwarding Strategies in Wireless Ad Hoc Networks – The Static Case. Tech. Rep. IC/2003/33, Swiss Federal Institute of Technology (2003)
5. Fudenberg, D., Tirole, J.: Game Theory. MIT Press, Cambridge Mass. (1991)
6. Goldsmith, A.J., Wicker, S.B.: Design Challenges for Energy-Constrained Ad Hoc Wireless Networks. IEEE Wireless Communications 4 (2002) 8-27
7. Kamoun, F.: A Drop-and-Throttle Flow Control Policy for Computer Networks. IEEE Trans. Comm. COM-29 (1981) 444-452
8. Konorski, J., Fair Packet Forwarding in Mobile Ad Hoc Networks: A Game-Theoretic Approach. Tech. Rep., Gdansk Univ. Tech. (2004)
9. Marti, S., Giuli, T. J., Lai, K., Baker, M.: Mitigating Routing Misbehavior in Mobile Ad Hoc Networks. In: Proc. 6th Annual Conf. on Mobile Computing and Networking MobiCom 2000 (2000) 255-265
10. Michiardi, P., Molva, R.: Making Greed Work in Mobile Ad Hoc Networks. Res. Rep. RR-02-069, Institut Eurecom, Sophia-Antipolis, France (2002)
11. Obreiter, P., Koenig-Ries, B., Klein, M.: Stimulating Cooperative Behaviour of Autonomous Devices – An Analysis of Requirements and Existing Approaches. Tech. Rep. 2003-1, Univ. of Karlsruhe, Germany (2003)
12. Srinivasan, V., Nuggehalli, P., Chiasserini, C. F., Rao, R. R.: Cooperation in Wireless Ad Hoc Networks. In: Proc. IEEE INFOCOM 2003 (2003)
13. Stallings, W.: Cryptography and Network Security: Principles and Practice. Prentice-Hall, Englewood Cliffs NJ (1999)
14. Urpi, A., Bonuccelli, M., Giordano S.: Modelling Cooperation in Mobile Ad Hoc Networks: a Formal Description of Selfishness. In: Proc. WiOpt 2003 Workshop on Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (2003)
15. Zhong, S., Chen, J., Yang, Y. R.: Sprite: A Simple, Cheat-Proof, Credit-Based System for Mobile Ad-Hoc Networks. In: Proc. IEEE INFOCOM 2003 (2003)