

Measuring Burden and Routing Fairness in Pocket Switched Networks

Tekenate E. Amah¹, Maznah Kamat¹, Kamalrulnizam Abu Bakar¹,
Waldir Moreira^{2,3}, Antonio Oliveira-Jr⁴, Marcos A. Batista⁵

¹Department of Computer Science, Faculty of Computing, Universiti Teknologi
Malaysia (UTM), 81310 Skudai, Johor, Malaysia

²Fraunhofer-AICOS - Porto, Portugal

³PPGMO, Federal University of Goiás - Catalão, Goiás, Brazil

⁴Institute of Informatics, Federal University of Goiás - Goiânia, Goiás, Brazil

⁵IBiotec, Federal University of Goiás, Catalão, Goiás, Brazil

amatek008@yahoo.com, {kmaznah, knizam}@utm.my,
waldir.junior@fraunhofer.pt, antonio@inf.ufg.br,
marcos.batista@pq.cnpq.br

Abstract. *A Pocket Switched Network (PSN) is formed by users carrying portable handheld devices such as smartphones and tablets, which store messages, carry them from one point to another via physical movement, and forwards them when a communication opportunity arises. The success of the network thereby depends on the willingness of users to participate. PSN protocols tend to subject most of the routing burden on only a smaller set of popular nodes. This results in drastic resource consumption on popular nodes, and may eventually lead to user dissatisfaction, withdrawal, and performance degradation of the network. The key to ensuring fairness in PSN routing lies in the ability to estimate the burden on nodes, utilize this knowledge to provide an acceptably fair utilization of node resources, and evaluate the level of fairness achieved. This paper is concerned with measuring: (i) the burden routing impacts on nodes; and (ii) the fairness of routing algorithms based on the distribution of this burden. First, we propose a Global Relative Burden Detection (GReBurD) mechanism to estimate the burden on nodes. Simulation experiments show that GReBurD is non-scenario specific and better infers the actual burden on nodes as compared with existing approaches. Next, we propose a new metric for evaluating the fairness of PSN forwarding algorithms that gives a better interpretation of the level of fairness implied.*

1. Introduction

Pocket Switched Networks (PSN) are based on the store-carry-forward (SCF) communication paradigm, in which messages are stored in node memory, physically carried from one point to another via user movement, and forwarded (through an available wireless communication interface such as Bluetooth or Wi-Fi) to another node when a communication opportunity arises. With respect to SCF-based networks, Mtibaa and Harras (2013) define fairness as the relative equality in the distribution of resource usage among neighbouring nodes in the network. Therefore, a forwarding algorithm is

absolutely fair if it subjects equal burden (i.e., resource expenditure) on all nodes. Absolute fairness leads to undesirable performance degradations in PSNs. This is due to the small world nature of the network (Orlinski and Filer, 2012), in which a relatively small set of “popular nodes” have more encounter opportunities than others (Tsvetovat and Kouznetsov, 2011; Muchnik *et al.*, 2013), thereby causing large variances in forwarding ability. Such nodes are often carried by popular users (e.g., politicians) or by those who travel around more frequently than others (e.g., deliverymen).

Nevertheless, it is important to achieve an acceptable level of routing fairness in order to ensure collaboration in PSNs. Since user devices are fairly homogeneous (Pujol *et al.*, 2009) in terms of resources such as battery capacity and computational capability¹, subjecting only few nodes to most of the burden in the network results in drastic resource consumption on them, and may eventually lead to user dissatisfaction, withdrawal, and performance degradation of the network (Amah *et al.*, 2016). Unfortunately, most routing protocols operate in this manner because their forwarding techniques and utilities are biased towards popular nodes (Pujol *et al.*, 2009). Considering that nodes are owned by humans, who may be unwilling to allocate most of their resources to routing, ensuring user participation is paramount. Hence, the success of routing solutions in real-world implementation lies in their ability to guarantee an acceptable level of fairness by sparingly utilizing the resources available on each node, instead of subjecting only few nodes to most of the burden in the network.

Besides the small world nature of PSNs, the distributed nature of the network makes it even more challenging to achieve fairness. Network designers need to be able to estimate the burden on nodes, utilize this knowledge to provide an acceptably fair utilization of node resources without significantly compromising network performance, and evaluate the level of fairness achieved. Existing approaches infer the burden on nodes from buffer information. However, inferring the burden on nodes from buffer information limits existing solutions from guaranteeing actual fairness in realistic scenarios, especially when other resource constraints (e.g., energy and processing) are considered. Focusing on the buffer alone may not achieve the desired level of fairness, because there is no guarantee that nodes will have equal buffer sizes. Furthermore, buffer occupancy is not a good indicator of how much burden a networking node is subjected to, due to the fact that node resources could be overwhelmed even without having a significant portion of the buffer occupied.

In this regard, our contribution is twofold. With respect to the first contribution (cf., Section 2), we analyse existing techniques for estimating the burden on nodes and identify their major drawbacks. We then propose the Global Relative Burden Detection (GReBurD) mechanism for estimating burden in PSNs without the identified drawbacks. Particularly, simulation experiments show that GReBurD is non-scenario specific and better infers the actual burden on a node, which is in accordance with energy consumption. With respect to the second contribution (cf., Section 3), we analyse existing metrics for measuring fairness and point out shortcomings of using these metrics to evaluate PSNs. We then propose a metric that is free of the identified shortcomings when used to measure routing fairness in PSNs. Our proposal retains all the properties of a desirable fairness metric, and allows for a better interpretation of the

¹ Note that allocated storage space may be according to user discretion, thus, may vary across nodes.

level of fairness implied. Finally, we present our conclusions and future work in Section 4.

2. Measuring Node Burden in PSNs

In this section, we analyse existing techniques for estimating the burden on nodes, identify major drawbacks in PSNs, propose an improved measure for burden, and support our claims through simulation experiments.

2.1. Existing measures for node burden

Authors often infer the burden on nodes from buffer information in order to achieve fair routing in PSNs. Pujol *et al.* (2009) forward messages to relay nodes based on the size of their message queue in order to balance load among nodes in the network. Grundy and Radenkovic's (2010) forwarding approach prioritizes nodes that have a higher percentage of remaining storage capacity, in order to distribute load away from popular nodes to less popular nodes. Mtibaa and Harras (2013) infer the burden on nodes from the number of messages they can carry, in order to ensure an efficiency fairness trade-off in forwarding. In order to achieve a tuneable trade-off between efficiency and fairness, Akestoridis *et al.*'s (2014) criterion for accepting a message is based on a utility derived from buffer occupancy.

Unfortunately, there is no guarantee that users in real-life will allocate fairly equal amount of storage space to routing. Therefore, comparing nodes by their buffer occupancy may be unfair in certain scenarios. For instance, depending on the amount of storage space allocated to routing, a node with 50% buffer occupancy may have received more messages and thereby expended more energy than another node with 75% buffer occupancy. By comparing buffer occupancies, messages continue to be directed towards the node with 50% buffer occupancy instead. This leaves a lot of room for error in real-world implementation, considering allocated storage resources may not be homogeneous after all.

This reveals that Akestoridis *et al.*'s (2014) approach will subject more burden on nodes that have assigned more storage space to routing since burden is considered as the ratio between remaining storage space and total storage capacity. Grundy and Radenkovic (2010) make forwarding decisions based on the percentage of available buffer space such that messages are directed towards nodes with less buffer occupancy. In a scenario where some nodes have allocated more storage space to routing than others, the former will be subjected to unfair treatment (e.g., more energy will be consumed from them), which cannot be detected by buffer occupancy.

Apart from the issue of different buffer sizes, buffer occupancy is still not a good indicator of how much burden a node is subjected to, since other resources could be overwhelmed without occupying a significant portion of the buffer. For instance, if the rate of receiving and sending messages is high, utilizing a fixed buffer occupancy threshold to ensure fairness may not be suitable. In that case, popular nodes may have used up all their energy allocated to routing without even reaching the threshold. Moreover, unlike the Internet, PSNs are constrained in terms of resources allocated to routing. Therefore, Internet-based approaches may not be suitable: even before the issue reflects on the buffer, energy usage may have exceeded the allocated quota. Unfortunately, most related work is oblivious of energy. Hence, the node continues to

function in the network even after allocated energy would have been exceeded, and the actual burden on popular nodes is not detected early enough.

Due to differences in the rate of sending and receiving messages, a higher buffer occupancy may not always mean higher burden, and vice versa. It is likely for popular nodes to have higher buffer occupancy most of the time, since they receive more messages. However, this may not always be the case, as popular nodes are also able to deliver more messages than less popular nodes due to higher encounter opportunities. As a result, depending on the rate of message generation, popular nodes may free their buffers faster than less popular nodes, especially when there is a drop in data traffic (e.g., after traffic bursts).

Detecting the burden on nodes via energy could also come to mind². However, unless the portion of energy allocated to routing can be monitored, overall energy consumption (i.e., residual energy) itself is also not a good indicator of the burden on a node. The chances that a device is low on battery due to other applications besides routing cannot be ruled out in real-life. This implies that determining the burden routing has impacted upon a node from residual energy may not be fair. For instance, a user conserving his battery for later use may find that most of it has been utilized for routing instead, while less would have been utilized if he was less conserving. Hence, inferring fairness from buffer occupancy or residual energy is only suitable for evaluation in controlled environments such as experimental simulations. To prevent bias in performance evaluation, simulation environments can be controlled to assume that device resources are consumed through participation in the network alone. Likewise, equal buffer size could be assigned to every node. This is however, not true in real-life, and should not be used as a basis of designing burden measures intended for real-world implementation. To address this gap, we propose an effective mechanism for inferring the burden on PSN nodes in Section 2.2.

2.2. Proposed measure for node burden

The approach employed here, namely Global Relative Burden Detection (GReBurD), determines the burden on a node from transmitted (i.e., received or sent) messages. For each node a , GReBurD maintains a counter for transmissions (i.e., received messages, excluding messages for which a is the destination, plus sent messages, excluding messages for which a is the source), $T(a)$. At the end of every time slot, $\Delta T b_i$ ($i \in [1, \infty]$), the current burden on a , $B_i(a)$, is computed by Equation 1, where $B_{i-1}(a)$ represents the burden on a in the previous slot, $\Delta T b_{i-1}$, if there was one (i.e., if $i > 1$). After computing the burden, $T(a)$ is reset to 0 for the next time slot, $\Delta T b_{i+1}$. Without resetting the transmission counter, the burden on nodes does not change even after a long period of inactivity.

$$B_i(a) = \begin{cases} T(a), & i = 1 \\ (T(a) + B_{i-1}(a))/2, & i > 1 \end{cases} \quad \text{Equation 1}$$

The burden at a given instant may be used to make forwarding decisions. In that case, the instantaneous burden can also be computed with Equation 1 (note that the

² Note that energy-aware routing solutions consider energy level in order to improve routing performance, rather than (achieving fairness by) making routing decisions according to the burden routing itself has impacted on nodes.

counter is not reset until the current time slot is elapsed). Based on the proposed mechanism, the instantaneous burden on nodes increase as they continue to participate in forwarding. At the beginning of a new time slot, the instantaneous burden on nodes is reduced since the transmission counter is reset to 0. Hence, the choice of $\Delta T b_i$ reflects the restoration experienced during the rest periods, such as when a device could be recharged, as opposed to periods of high traffic such as rush hours. For this, we have selected a period of six hours and leave details of this to future work. This choice is based on results from several trials, where we found that it is likely to reset the transmission counters when the user is either at home or at work, i.e., when there are least transmission activities in the network.

2.3. Evaluation of proposed measure for node burden

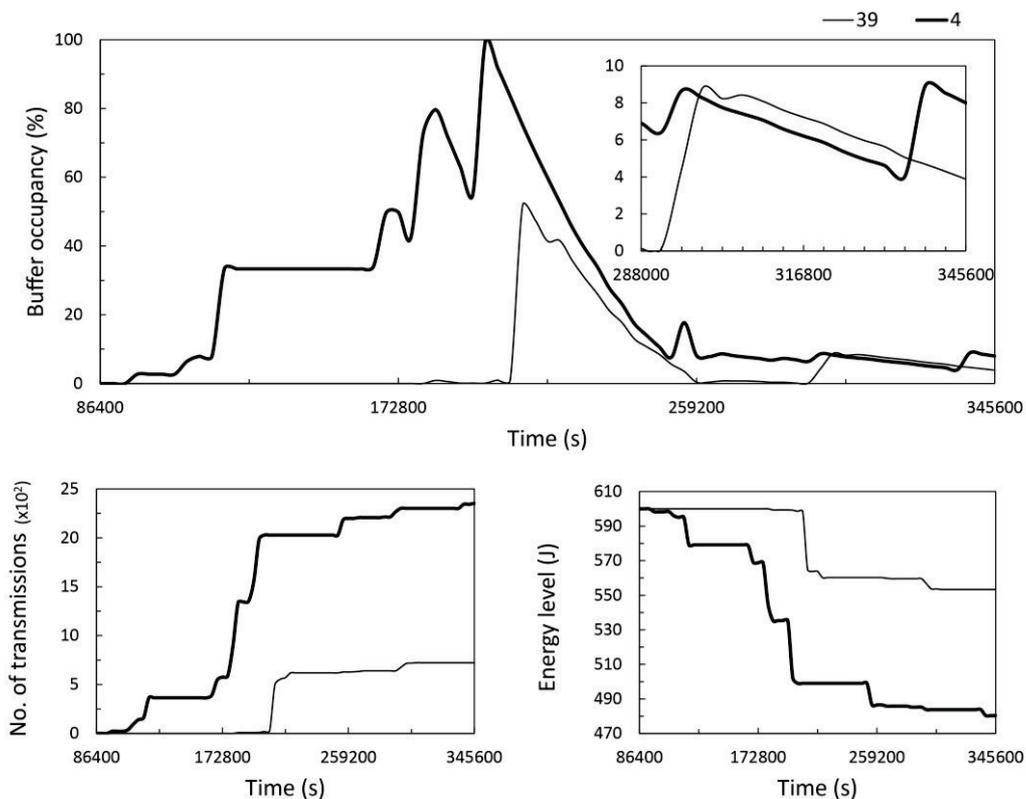
Our evaluation was carried out in the Opportunistic Network Environment (ONE) simulator (Keränen *et al.*, 2009). We assume a scenario in which portable handheld devices, carried by mobile users in a city, convey data between static sources (e.g., sensors) and destinations (e.g., base stations) placed in strategic locations. Sources are situated around homes while destinations are located around bus stops and meeting spots in the Helsinki scenario in ONE simulator. Nodes forward messages according to the PROPHET routing protocol (Grasic *et al.*, 2011). For evaluation purposes, all nodes are given equal resources in terms of buffer size and initial energy. Since our interest is only on energy consumed due to message forwarding, we set scan energy (i.e., energy consumed from device discovery), scan response energy (i.e., energy consumed from device discovery response), and base energy (i.e., energy consumed in idle state) to 0, in order to avoid bias. All mobile nodes move according to the Working Day Movement model (Ekman *et al.*, 2008), which presents realistic human movement patterns. The parameters used for the simulation setup are summarised in Table 1.

Table 1. Simulation parameters.

Parameter	Value	Parameter	Value
Wireless interface	Bluetooth	Receive/transmit energy	0.08 mW/s
Transmission rate	2 Mbps	No. of source nodes	80
Buffer size	10 Mb	No. of destination nodes	36
Message size (uniformly distributed)	10 Kb to 15 Kb	No. of buses	18
Simulation area	7 km × 8.5 km	No. of taxis	50
Initial energy	600 J	No. of pedestrians	416

We have chosen a simulation duration of 4 days. 1 day is allowed before message generation, for warm-up. The warm-up period allows nodes to acquire enough routing information to make forwarding decisions according to the routing protocol in use – in this case PROPHET. Data is generated for three days: 1st day, high generation rate (1msg/hr per src); 2nd day and 3rd day, low generation rate (1msg/24hrs per src). This allows us to investigate the performance of the burden measures under changing data traffic rates, as previously mentioned (cf., Section 2.1). Each message is assigned a TTL of 1 day, so that node buffers are freed when messages are delivered to their destinations, dropped (according to the FIFO policy) due to buffer overflow, or as a result of TTL exhaustion. For our proposed GReBurD, we have obtained results with a script we have implemented in ONE simulator. With the script, we are able to obtain the burden on each node at different points in time.

We address the following question: "is there a scenario that could render buffer occupancy inaccurate for determining the relative burden on nodes due to changes in the rate of message generation?" By comparing the results of randomly selected node pairs, some pairs verify that buffer occupancy may not always be an accurate representation of the burden on nodes. In fact, instantaneous buffer occupancy depends on various variables that may not correlate with the relative amount of contribution a node has done in forwarding messages, e.g., TTL, the number of destinations that can be directly encountered, the queuing policy in use, and routing protocol conditions for dropping messages. Figure 1a, 1b and 1c compares three node pairs, Node 39 and 4, Node 41 and 0, and Node 48 and 5, respectively. During the three days of message generation, the second node of each pair (i.e., Node 4, 0 and 5) – which we term as the popular node – transmits more messages (i.e., the sum of sent and received messages), and this is reflected in their energy level. Hence, the less popular node, carrying more messages in its buffer at some point in time does not make its burden exceed that on the popular node, unless this continues for a period of time that is able to compensate for the burden that the popular node has been subjected to. In addition, higher buffer occupancy on the less popular node would mean higher burden only when it has experienced more transmissions – note that at this point in time, it is possible for the node with lower buffer occupancy to be delivering more messages to destinations, hence, freeing its buffer at a faster rate. The figures show that the buffer occupancy on the less popular node exceeds that on the popular node for considerable hours (11, 2 and 12 hours for Figure 1a, 1b and 1c, respectively), during which false positives – in terms of burden – could arise if routing decisions are made solely based on buffer occupancy.



(a)

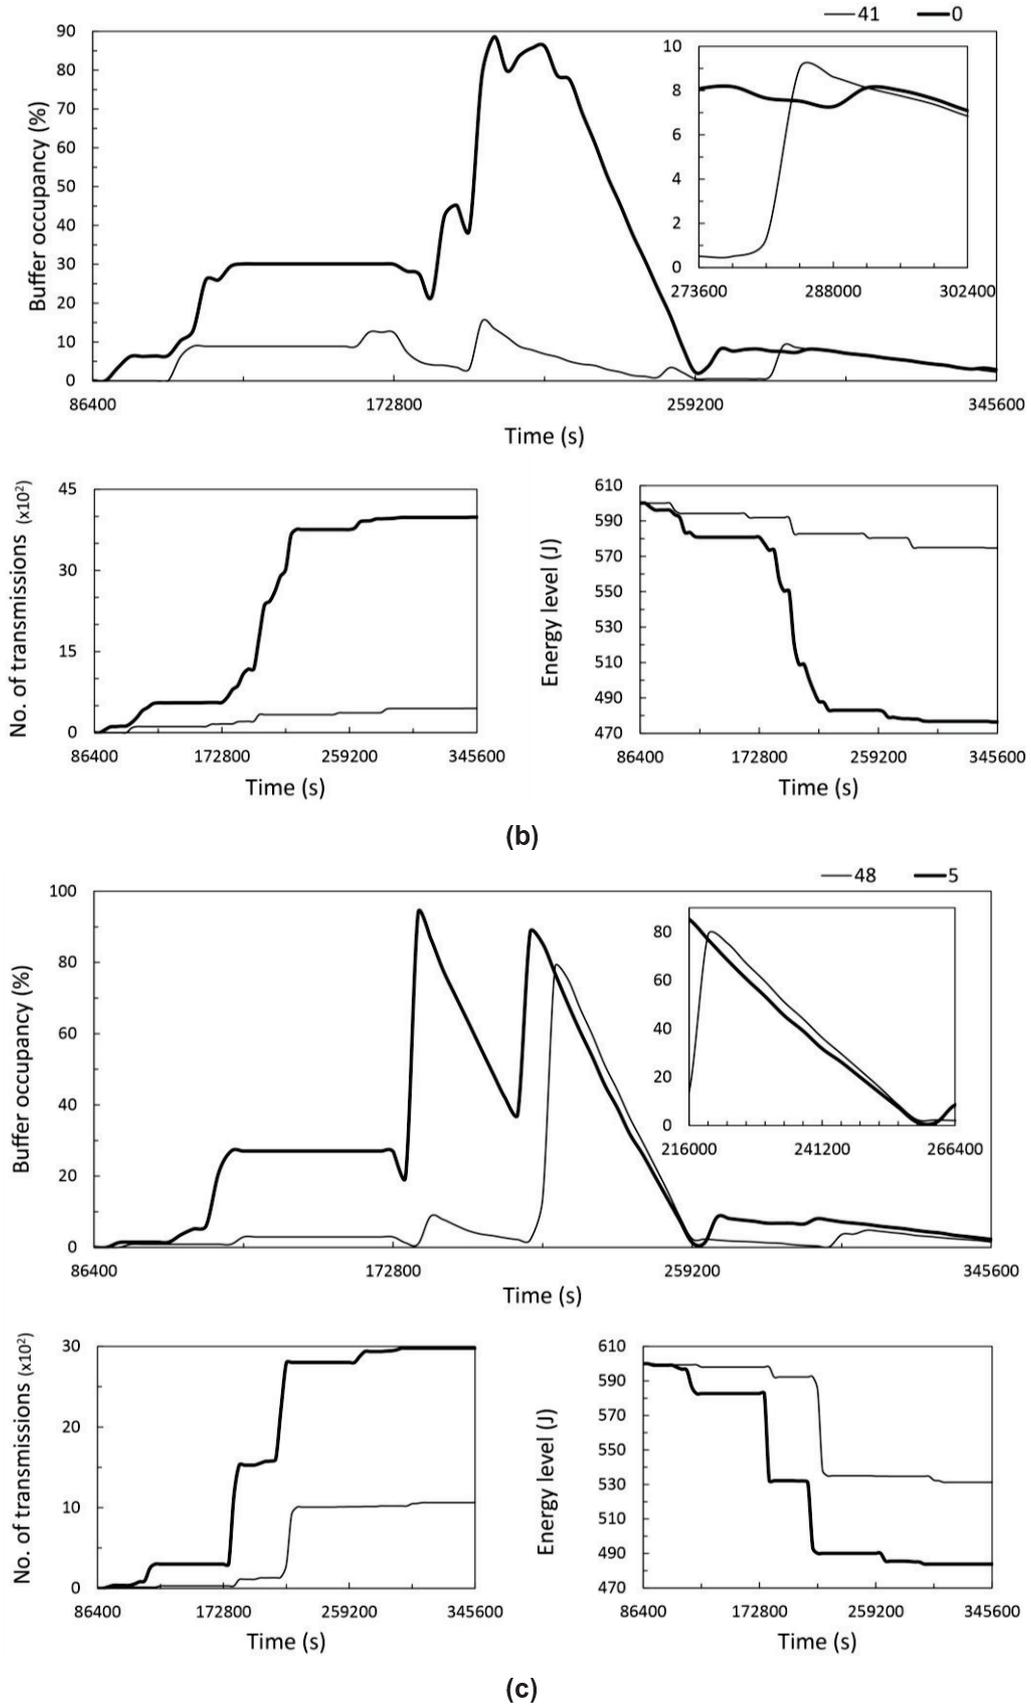


Figure 1. Buffer occupancy, number of transmission and energy level on: (a) Node 39 and 4 (b) Node 41 and 0 (c) Node 48 and 5.

The results in Figure 1 have shown that the instantaneous buffer occupancy may not be able to reflect the amount of transmissions a node has experienced, or the energy it has expended in doing so. By comparing buffer occupancies in the highlighted sections of Figure 1, messages would be directed away from the less popular node and towards the popular node at some point in time. Such approaches try to balance the buffer occupancy of nodes, hence, achieve fairness in terms of buffer occupancy. However, they fail to achieve fairness in terms of the energy consumed in routing, since instantaneous buffer occupancy does not always indicate the actual burden on a node. By observing Figure 1, the current state of the buffer and snapshots of previous states could be used to compute the instantaneous burden. However, there is no guarantee that the snapshot of a previous state would be captured at point in time when the buffer would be able to reflect the actual burden on the node.

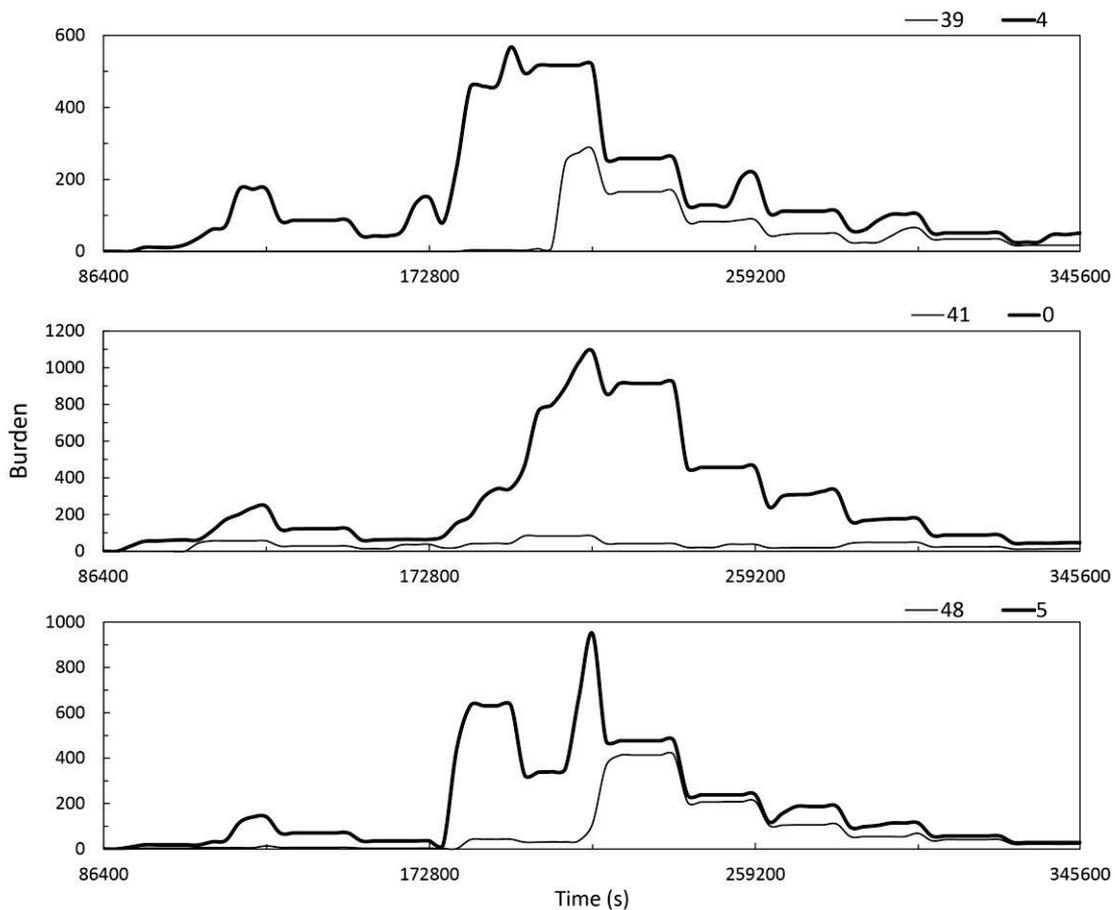


Figure 2. The burden on the nodes according to GReBurD.

The results from this experiment makes it easier to imagine how inferring node burden from buffer occupancy would be a problem in a scenario where every user's device is a potential destination. Popular nodes would then be able to encounter more destinations directly and free up their buffers – lecturers and their students, the bus driver's device and encounters with passenger devices, sales people who serve a lot of customers, etc. Without a steady flow of incoming messages, the buffer occupancy on such nodes may drop at a relatively higher rate, thereby rendering buffer occupancy momentarily inaccurate for determining the relative burden on nodes. To address this issue, our approach considers the number of transmissions within a time frame, and tries

to account for any periods of inactivity between. In other words, our proposed burden measure decreases with reducing number of transmissions and vice versa. Therefore, the burden on a less popular node in Figure 1 will surpass that on its popular counterpart only when the former has done more transmissions and expended more energy. As opposed to inferring the burden from the number of transmissions alone, the reduction in burden when there is less number of transmissions is able to account for periods of inactivity, during which device resources could be replenished (e.g., battery recharge). We believe that our proposed GReBurD is also more suitable in real-life scenarios because it can cope with the dynamicity of node behaviour. Suppose Node A just joined the network. By inferring burden from the number of transmissions alone, it may take a very long time for the burden on Node A to reach that on the other nodes, even if it is a popular node. With our proposed burden measure on the other hand, Node A would be able to catch up in the next counter reset interval. As shown in Figure 1, the burden on the less popular nodes never surpasses that on their popular counterpart, even when the buffer occupancy of the former exceeds that of the latter. This burden corresponds to the number of transmissions and energy expended on both nodes as shown in the figure.

3. Measuring Routing Fairness in PSNs

In this section, we analyse the metrics for fairness, point out shortcomings of using these metrics to evaluate PSNs, and propose a metric specifically for measuring routing fairness in PSNs.

3.1. Existing metrics for routing fairness

In PSNs, the fairness metric measures “evenness” of burden distribution on nodes, and if not even, indicates how far the distribution is from evenness. According to Jain *et al.* (1984), the index should have the following properties in order to give an intuitive understanding of fairness: (i) independent of population size, applicable to any number of nodes; (ii) independent of scale and metric, should give the same results across different units of measurement; (iii) bounded between 0 and 1, a totally fair and a totally unfair distribution should have an index of 1 and 0, respectively; and (iv) continuous, ability to reflect any slight change in distribution.

Various measures are used to determine the level of fairness among nodes in homogeneous networks. Pujol *et al.* (2009) measure fairness by the fraction of nodes that carry out 50% of the total number of forwards in the network. Thus, a larger fraction implies more fairness and vice versa. The metric used by Soelistijanto and Howarth (2012) is the ratio of maximum to mean data traffic seen by nodes in the network, with a lower value implying a more even distribution of traffic among the nodes. In Mashhadi *et al.*'s (2012) evaluation, fairness is given by the coefficient of variation of the total load forwarded by nodes. The index proposed by Mtibaa and Harras (2013) is the difference between the message load distribution given by the forwarding process and the uniform distribution among nodes.

However, none of these measures possesses all the aforementioned desired properties of a suitable fairness index. For instance: continuity, the drawback of the measures used by Pujol *et al.* (2009), and Soelistijanto and Howarth (2012) is similar to that of the max-min ratio (Marson and Gerla, 1982), as they do not reflect slight changes in distribution; and boundedness, the index proposed by Mtibaa and Harras (2013) ranges from positive to negative values and the metric used by Mashhadi *et al.*

(2012) drifts towards $-\infty$ as fairness decreases. As a result, it is not easy to interpret the level of fairness implied by the existing measures. In order to meet account for these criteria, previous work (e.g., Akestoridis *et al.*, 2014; Fan *et al.*, 2014) use Jain *et al.*'s (1984) index (cf., Equation 2) to measure fairness in homogeneous networks.

$$f_A(x) = \frac{[\sum_{i=1}^n x_i]^2}{n \sum_{i=1}^n x_i^2}, \quad x_i \geq 0 \quad \text{Equation 2}$$

In Equation 2, $f_A(x)$ is the fairness index of a protocol A , that distributes x amount of burden to n nodes, such that the i^{th} node receives a burden x_i .

Table 2. Fairness of distributing a total burden of 10 to two nodes in a homogeneous network according to Jain *et al.* (1984).

	Algorithm and burden distribution					
	A	B	C	D	E	F
Node 1	10	9	8	7	6	5
Node 2	0	1	2	3	4	5
Fairness (%)	50	60.98	73.53	86.21	96.15	100

From Table 2, algorithm A is totally unfair since node 1 bears the entire burden in the network and F is totally fair since the total burden is equally distributed between both nodes. Thus, $f_A(x)$ is expected to be 0 and $f_F(x)$ is expected to be 1, which corresponds to a fairness of 0% and 100%, respectively. However, Jain *et al.*'s (1984) index assigns a fairness of 50% to algorithm A – instead of 0% – which is not a suitable interpretation of fairness in PSNs. According to Jain *et al.* (1984), “a distribution algorithm with a fairness of 0.1 means that it is unfair to 90% of the users”. This implies that algorithm A is unfair to only 50% of the nodes, hence, results in a 50% fairness. In terms of evaluating fairness in a PSN, however, this scenario would be better interpreted as total unfairness, which means a 0% fairness. To account for this, we propose a new metric for measuring PSN routing fairness in Section 3.2.

3.2. Proposed metric for routing fairness

To illustrate our proposed fairness metric, consider a scenario of six fictitious forwarding algorithms, A to F, with 4 nodes each, n1 to n4, that are burdened at different extents for each forwarding algorithm (cf., Table 3). The standard deviation from the mean burden for a forwarding algorithm, a , is given by Equation 3.

$$\beta_a = \sqrt{\sum_{i=1}^n (b_{a,i} - B_a)^2 / n} \quad \text{Equation 3}$$

Where n is the total number of nodes, $b_{a,i}$ is the burden experienced by node i running on algorithm a , and B_a is the mean burden experienced by a node under algorithm a . Then $\beta_{a,max}$, the maximum possible standard deviation for algorithm a (i.e., the standard deviation if only one node bears all the burden under algorithm a , which is also the least fair scenario, such as algorithm F in Table 3), is given by Equation 4.

$$\beta_{a,max} = \sqrt{\left[(b_{a,total} - B_a)^2 + B_a^2(n - 1) \right]} / n \quad \text{Equation 4}$$

Where $b_{a,total}$ is the total burden, i.e., sum of the burdens experienced by every node in the network. As given by Equation 5, the ratio between β_a and $\beta_{a,max}$ indicates the unfairness of algorithm a .

$$U_a = \frac{\beta_a}{\beta_{a,max}} \times 100 \quad \text{Equation 5}$$

$$F_a = 100 - U_a \quad \text{Equation 6}$$

If every node experiences equal burden for algorithm a , F_a in Equation 6 will become 100% (e.g., A in Table 3). Likewise, if only one node bears all the burden, F_a will become 0% (e.g., F in Table 3). F_a also decreases accordingly, the less evenly the total burden is distributed among the nodes (e.g., B to E in Table 3). Therefore, F_a can be used to measure the fairness of a forwarding algorithm, a , in terms of the burden on nodes.

Table 3. Six fictitious forwarding algorithms for describing the working principle of the proposed fairness metric and how it interprets fairness.

Forwarding algorithm (a)	Nodes and burden				β_a	Fairness; F_a (%)
	n1	n2	n3	n4		
A	12	12	12	12	0	100
B	13	12	12	11	0.7	96.6
C	24	8	8	8	6.9	66.7
D	32	8	4	4	11.7	43.9
E	47	1	0	0	20.2	2.8
F	48	0	0	0	20.8	0

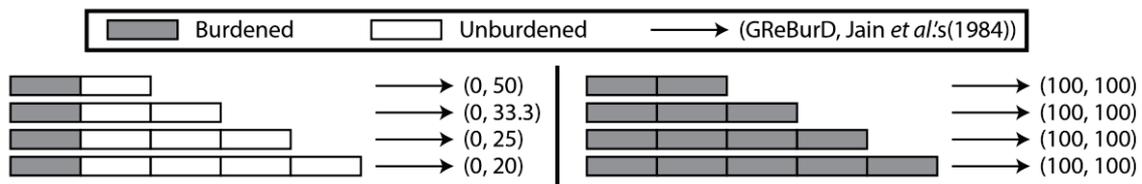
3.3. Evaluation of proposed metric for routing fairness

Our proposed fairness metric is applicable to any number of nodes, gives the same results across different units of measurement, results in 0 and 1 for a totally fair and a totally unfair distribution, and has the ability to reflect any slight change in distribution. Here, we compare the performance of our proposed fairness metric with Jain *et al.*'s (1984), which is the only one among the existing metrics that possess all the stated desirable properties of a fairness metric. As shown in Table 4, we allocate 10 items to two nodes according to 6 different allocation schemes. In a PSN scenario, routing fairness should be 100% when the two nodes experience equal burden, and both fairness metrics are able to reflect this for Scheme 6. On the other hand, burdening only one node in a PSN should result to a 0% fairness, which we expect for Scheme 1. With Jain *et al.*'s (1984) metric, Scheme 1, the most unfair allocation scheme, gives a fairness of 50%, while our proposed metric gives a fairness of 0%. With this behaviour, it is easier to give a better interpretation of the level of fairness implied by our proposed fairness metric. As shown in Table 4, our proposed fairness metric is able to reflect a totally unfair and a completely fair allocation of burden in PSNs. The fairness of a forwarding algorithm can be computed from the burden measure we proposed and evaluated in Section 2, using the instantaneous burden on each node.

Table 4. Fairness of different schemes for allocating 10 items to two nodes.

Allocation scheme	Allocated items		Fairness metric (%)	
	Node A	Node B	Proposed	Jain <i>et al.</i> 's
1	10	0	0	50.0
2	9	1	20	61.0
3	8	2	40	73.5
4	7	3	60	86.2
5	6	4	80	96.2
6	5	5	100	100.0

The difference between the two metrics is in their interpretation of minimum fairness. First, Jain *et al.*'s (1984) metric considers the entire population. Hence, the extent of fairness interpreted when only one node bears all the burden changes with the total population (cf., Figure 3). Although reasonable, fairness never reaches 0%, and one is forced to bear in mind the total population in order to fully understand the interpretation – since the most unfair scenario could be represented by different values. Second, the condition for Jain *et al.*'s (1984) metric to give 0% fairness is when the numerator of Equation 2 equals 0. This only occurs when negative burden values are considered, i.e., 10 and -10 gives a 0% fairness in this case. However, since burden values are usually non-negative, Jain *et al.*'s (1984) metric may require them to undergo specific normalization processes in order to be bounded between 0 and 1 (e.g., representing 10 and 0 by 5 and -5, respectively).

**Figure 3. Fairness according to our proposed metric and Jain *et al.*'s (1984).**

Considering that instantaneous burden on nodes changes over time, burden at different points in time may have to be averaged for each node in order to obtain the actual fairness of a routing algorithm. Using this method and a granularity of 3600 seconds, we are able to compute the fairness of the scenario in Section 2.3 with our proposed metric and Jain *et al.*'s (1984) metric as 95.44% and 44.54%, respectively.

4. Conclusions and Future Work

In a Pocket Switched Network (PSN), portable handheld devices, such as smartphones and tablets, store messages, carry them through user movement, and forward them to other devices when a communication opportunity arises. The network thus, depends on the willingness of users to participate and share their devices as routers. Due to the versatility of portable handheld devices, the availability of required resources is limited. Furthermore, users may not be willing to shed all their resources on behalf of the network. Therefore, it is important that PSN protocols be as fair as possible by sparingly utilizing the resources available on each node instead of subjecting most of the routing burden to only a few set of popular nodes. Overlooking fairness results in drastic resource consumption on popular nodes, and may eventually lead to user dissatisfaction, withdrawal, and performance degradation of the network.

In order to ensure fairness in PSN routing, network designers need to be able to estimate the burden that has been impacted on nodes, utilize this knowledge to provide an acceptably fair utilization of node resources, and evaluate the level of fairness achieved. Existing approaches for ensuring routing fairness rely on buffer occupancy to indicate the level of burden on nodes. Unfortunately, as experiments in Section 2 have revealed, buffer occupancy may not always be able to reflect the amount of burden that a node has been subjected to. In this regard, we have proposed GReBurD, a mechanism that estimates the burden on PSN nodes, and evaluated its performance in ONE simulator. GReBurD is simple yet effective, as it is non-scenario specific and better infers the actual burden on a node as compared with existing approaches.

We also analysed existing metrics for measuring fairness in the distribution of burden among nodes. We identified that most of the metrics fail to satisfy the desirable properties of a fairness metric, while the most accepted metric leaves room for improvement in terms of interpreting the level of fairness implied. In this regard, we have proposed a new metric with which network designers can evaluate the fairness of PSN forwarding algorithms. The proposed metric possesses all the desirable properties of a fairness metric, and also gives a better interpretation of the level of fairness implied in PSNs.

Future work includes a distributed version of GReBurD that can locally estimate the relative burden on a node, without requiring synchronization between devices. Future work also includes investigating the impact of different reset intervals on the performance of GReBurD. For the future version, we also consider a means of resetting the transmission counter at optimum periods detected from device activity history, instead of selecting fixed intervals. We also have planned to evaluate the GReBurD in a real-life scenario.

Acknowledgment

Thanks are due to Fundação de Amparo à Pesquisa do Estado de Goiás (FAPEG) and PNP/DCAPES.

References

- Akestoridis, D., Papanikos, N., and Papapetrou, E. (2014). "Exploiting social preferences for congestion control in opportunistic networks," 2014 IEEE 10th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob), Larnaca, pp. 413-418. DOI: 10.1109/WiMOB.2014.6962204.
- Amah, T., Kamat, M., Moreira, W., Abu Bakar, K., Mandala, S., and Batista, M. (2016). Towards next-generation routing protocols for pocket switched networks, *Journal of Network and Computer Applications*, Volume 70, pp. 51–88. DOI: 10.1016/j.jnca.2016.05.011.
- Ekman, F., Keränen, A., Karvo, J., and Ott, J. (2008). Working Day Movement Model. *MobilityModels'08*, Hong Kong SAR, China.
- Fan, X., Li, V., and Xu, K. (2014). "Fairness Analysis of Routing in Opportunistic Mobile Networks," In *IEEE Transactions on Vehicular Technology*, vol. 63, no. 3, pp. 1282-1295. DOI: 10.1109/TVT.2013.2282341.

- Grasic, S., Davies, E., Lindgren, A., and Doria, A. (2011). The evolution of a DTN routing protocol—prophetv2. In: Proceedings of the 6th ACM workshop on challenged networks, CHANTS '11. ACM, New York, pp 27–30. DOI:10.1145/2030652.2030661.
- Grundy, A., and Radenkovic, M. (2010). "Promoting congestion control in opportunistic networks," 2010 IEEE 6th International Conference on Wireless and Mobile Computing, Networking and Communications, Niagara Falls, ON, pp. 324-330. DOI: 10.1109/WIMOB.2010.5645048.
- Jain, R., Chiu, D., and Hawe, W. (1984). "A quantitative measure of fairness and discrimination for resource allocation in shared computer systems," Digital Equipment Corporation, Tech. Rep. DEC-TR-301.
- Keränen, A., Ott, J., and Kärkkäinen, T. (2009). The ONE simulator for DTN protocol evaluation. Simutools '09, Proceedings of the 2nd International Conference on Simulation Tools and Techniques. Rome, Italy. No. 55. DOI: 10.4108/ICST.SIMUTOOLS2009.5674.
- Marson, M., and Gerla, M. (1982). "Fairness in Local Computing Networks," In Proceeding of IEEE ICC.
- Mashhadi, A., Mokhtar, S., and Capra, L. (2012). Fair content dissemination in participatory DTNs, Ad Hoc Networks (Special Issue on Social-Based Routing in Mobile and Delay-Tolerant Networks), Volume 10, Issue 8, November 2012, Pages 1633–1645. DOI: 10.1016/j.adhoc.2011.05.010.
- Mtibaa, A., and Harras, K. (2013). Fairness-related challenges in mobile opportunistic networking, Computer Networks, Volume 57, Issue 1, pp. 228-242. DOI: 10.1016/j.comnet.2012.08.019.
- Muchnik, L., Pei, S., Parra, L., Reis, S., Andrade, J., Havlin, S., and Makse, H. (2013). Origins of power-law degree distribution in the heterogeneity of human activity in social networks. Scientific Reports 3: 1783. DOI: 10.1038/srep01783.
- Orlinski, M., and Filer, N. (2012). Quality distributed community formation for data delivery in pocket switched networks. In Proceedings of the Fourth Annual Workshop on Simplifying Complex Networks for Practitioners (SIMPLEX '12). ACM, New York, NY, USA, pp. 31-36. DOI: 10.1145/2184356.2184365.
- Pujol, J., Toledo, A., and Rodriguez, P. (2009). "Fair Routing in Delay Tolerant Networks," INFOCOM 2009, IEEE, Rio de Janeiro, pp. 837-845. DOI: 10.1109/INFCOM.2009.5061993.
- Soelistijanto, B., and Howarth, M. (2012). "Traffic distribution and network capacity analysis in social opportunistic networks," 2012 IEEE 8th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob), Barcelona, pp. 823-830. DOI: 10.1109/WiMOB.2012.6379171.
- Tsvetovat, M., and Kouznetsov, A. (2011). Centrality, Power, and Bottlenecks. Social Network Analysis for Startups. O'Reilly Media, Inc., 1005 Gravenstein Highway North, Sebastopol, CA 95472.