Tipping the Balance: Towards Fairer Energy Consumption in Groups Connected via Device-to-Device Technologies

Sascha Trifunovic, Andreea Picu, Theus Hossmann
Communication Systems Group
ETH Zurich, Switzerland
lastname@tik.ee.ethz.ch

ABSTRACT

By using device-to-device communication capabilities of smartphones, opportunistic networks promise to fill the gap of infrastructure-based networks in remote areas, to support communication in disaster and emergency situations, as well as to enable new local social networking applications. Yet, to become feasible in practice and accepted by the users, it is crucial that opportunistic communication is energy-efficient and that the load on mobile phone batteries is evenly or fairly distributed among users.

In this paper, we deal with the imbalance problem of today's device-to-device communication technologies (Wi-Fi Direct, Bluetooth): the "host" of a connection consumes two to five times more energy than a "client". For a *pair* of connected nodes, we have proposed a role switching scheme in previous work, in which we use the distribution of the remaining connection/contact time to derive an appropriate switching interval. Here, we lay the ground work for extending this solution to *groups* of more than two nodes. We first discuss several potential definitions of a group and its lifetime. Based on these definitions, we analyze the distribution of group lifetime in several real-world traces and show that, similarly to pairwise contact durations, group lifetimes obey a heavy-tailed distribution. Finally, we discuss how to use this generalization to design a role switching scheme for group connections, aimed at achieving a more balanced energy consumption among group members.

Categories and Subject Descriptors

C.2.1 [**Network Architecture and Design**]: Wireless Communications

Keywords

Opportunistic Networks, Energy, Fairness, Bluetooth, Wi-Fi Direct

1. INTRODUCTION

Today, most people carry mobile phones featuring technologies like Bluetooth or Wi-Fi Direct, which allow device-to-device communication. Using these technologies, two or more devices can exchange data whenever they are within mutual transmission range, thereby forming an opportunistic network [1, 2]. Such opportunistic networks were proposed as a solution to fill the gaps of the existing networking infrastructure in remote and rural areas [3], to enable communication when infrastructure breaks down during natural disasters [4], or to elude censorship. Moreover, opportunistic networks can mitigate the pressure on infrastructure, caused by exponentially growing traffic demands, by offloading some traffic [5].

State/Event	*Power/Energy*	*STD*
Wi-Fi Direct connected (station)	49.75 mW	3.90 mW
Wi-Fi Direct connected (AP)	231.92 mW	9.14 mW
Wi-Fi Direct connect (station)	3523.78 mJ	714.44 mJ
Wi-Fi Direct connect (AP)	1654.50 mJ	395.25 mJ

Table 1: Consumption of connection operations.

However, to make opportunistic networking feasible in practice, users must accept to contribute, despite their resource constrained phones. One critical challenge to achieving this is to minimize the impact of device-to-device networking on battery lifetime. In a recent study [6] of the energy consumption of Wi-Fi Direct, Wi-Fi and Bluetooth (most commonly used technologies in opportunistic networking), we showed that all of them are highly unbalanced in the formation and maintenance of peer-to-peer connections. More precisely, of two connected devices, the one that "hosts" the connection incurs a much higher energy cost than the "client" device. As shown in Table 1, the imbalance in energy consumption in Wi-Fi Direct is significant, as the host device consumes nearly five times as much as the client device. This is similar for other technologies, as well as for group connections of more than two devices.

An intuitive remedy for the above problem is to periodically alternate the hosting role among the participating devices; we refer to this as *role switching*. While the concept of role switching is straightforward, implementing it engenders a complex fairness–efficiency trade off. On the one hand, there is a non-negligible one-time cost associated with switching roles (i.e. "rebuilding" the connection(s)), which for the sake of efficiency imposes a minimum time to stay connected before switching. On the other hand, the time-frame of a connection opportunity is limited in opportunistic networks, which may require very frequent role switching to achieve a desired balance[1] in the energy consumption of individual devices. Thus, working out a good switching schedule to solve this trade off poses significant challenges.

In [6], we proposed a solution to this fairness–efficiency trade off, for the simple case of connections between two devices only and the goal of "equal share" fairness. Using the fact that the duration of pairwise connections (or *contacts*) in opportunistic networks follows a power law distribution [1], we derived a closed-form expression for the distribution of the remaining contact duration (de-

S3'13, September 30, 2013, Miami, Florida, USA

---

[1] This is determined by the chosen fairness definition.

	H06	ETH	SF
# pairwise contacts	128 971	22 958	1 339 273
# nodes	78	20	536
scanning interval	2 min	2 s	30 s

Table 2: Contact traces.

pending on the elapsed contact duration). Given this remaining contact duration, as well as the costs of switching and of maintaining the connection, we can easily calculate a fair and efficient role switching interval. Applying this to real connection traces shows a reduction in the energy required for role switching by up to 92%, compared to a static role switching interval, while maintaining a good level of fairness.

In this paper, we build the foundations for extending our role switching solution to group connections, where more than two devices are in contact with each other. To this end, we first define a group and its lifetime; we then analyze real-world contact traces from the perspective of group contacts (rather than pairwise), and find that the heavy tail property holds in this case as well. Similarly to the pairwise case, this implies that a node's remaining membership time in a group increases with the elapsed membership time. Based on this analysis, we explain how to adapt our role switching scheme, originally designed for pairs, to groups.

2. GROUPS AND THEIR LIFETIMES

Most studies of mobility-induced connection opportunities or contacts among nodes focus on pairs of nodes, under the assumption that the network is so sparse that the formation of larger groups is highly unlikely [1, 7]. They typically study inter-contact times, which are the most significant delay component in transporting data, and to a lesser extent contact durations. The few works that consider group contacts analyze the group formation process and the size distribution [8] or deal with routing over such groups [9]. There is, however, no work on precisely defining a group at a *microscopic* level, nor on group lifetime properties.

In opportunistic networks, nodes are likely unaware of any neighbors beyond the first hop, since connection opportunities are assumed to be short and must be primarily used for exchanging data, rather than control information. Thus, defining groups based solely on each node's local view is the most convenient approach. However, for the purpose of our role switching scheme, all the nodes of a group will need to collectively and periodically decide on who is/are the next host(s) for that group's connections. With this in mind, we adopt a global view of groups: *A set of nodes form of group at time t, if they all belong to the same connected component at that time.* To extract such groups and their dynamics from contact traces[2], we keep track of all available groups at each time point (event) of the trace. For illustration simplicity, we assume here that each group is a clique, such that only one host node is needed in any group. Relaxing this assumption is part of our ongoing work.

Based on this group definition, there are several possible ways to measure and define group lifetime. Since groups are highly dynamic in opportunistic networks, deciding when and how to reset the lifetime of a group is crucial, with major implications on the efficiency and fairness of role switching. In particular, as explained in the introduction, role switching only starts being useful after a minimum connection time $t_{\min}$ (to balance the one-time cost of switching), and has increasing chances of achieving fairness with increasing connection lifetime. We discuss below two possible group lifetime definitions, indicating their pros and cons.

Composition time. One of the simplest ways to define group lifetime is to restart time whenever the composition of the group changes (i.e. a node arrives or leaves). These times will naturally be shorter than pairwise contact times, as two nodes in contact will restart the time whenever a third node comes into contact. Short group lifetimes have obvious drawbacks for the purpose of role switching. To illustrate this, consider a long pairwise contact, often "perturbed" by short connections to some random third node. This results in short group lifetimes dictated by the short third node connections, with any memory of the long contact between the original two nodes erased at every new event. In this case, any role switching scheme will most of the time be waiting for the minimum connection time $t_{\min}$ to elapse, instead of trying to achieve long term even energy consumption by profiting from long connection times.

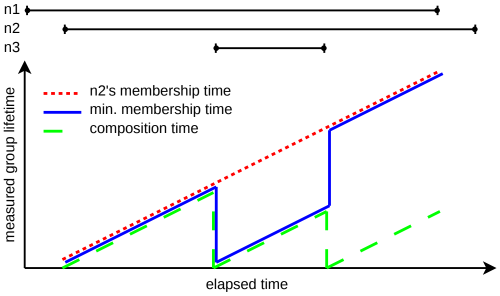

Figure 1: Group lifetime definition overview.

Minimal membership time. Let the *membership time* of a single node be the total time a node is member of a group (i.e. connected to any node until it is alone again). Then, the minimal membership time of a group is defined as the minimum over all the membership times in the group. Taking this as a measure of a group's lifetime is more suited to our purpose of role switching: it ensures that group lifetime is restarted when a new node joins (for the sake of fairness), but not when a node departure event restores the group to a previous state. In the above example, a role switching algorithm can pick up where it left off every time the short connection to a third node is broken.

Fig. 1 illustrates the above two definitions of group lifetime. The three horizontal lines at the top represent the time each of three nodes, n_1, n_2 and n_3, are present at some common location. Assuming they come in contact with each other when they are simultaneously present at the location, the graph underneath shows how group lifetime evolves over time, for each definition.

With clear definitions of groups and their lifetimes, we can now attempt to predict remaining group lifetime, based on the age of the group. Such a prediction is invaluable in the design of a good role switching scheme, as explained in the introduction. To do so, we analyze the distribution of group lifetime in real-world contact traces. We chose traces that are frequently used in opportunistic networking research: the Haggle 2006 trace (H06), collected during the three days of the Infocom conference in 2006 [7], and the ETH trace (ETH) collected on the ETH Zurich campus in 2005 [10]. We also use the San Francisco taxicab trace (SF), representing GPS position recordings for over 500 taxicabs over a period of a month [11]; the contacts in this trace are inferred based on a conservative transmission range of 30m. The characteristics of these datasets are summarized in Table 2.

For all these traces, we used the methods described in [12] (i.e. maximum likelihood (ML) for fitting, maximum likelihood ratio for comparison of fits, Kolmogorov-Smirnov for goodness-of-fit etc) to identify the most likely type of distribution for our group lifetimes, among the usual exponential, power law, truncated power law etc. We found that a Pareto distribution is the best fit for both the composition and the minimal membership time (Fig. 2).

---

[2] Typically recorded as a sequence of contact 4-tuples giving the two communicating nodes, a contact start time, and its end time.

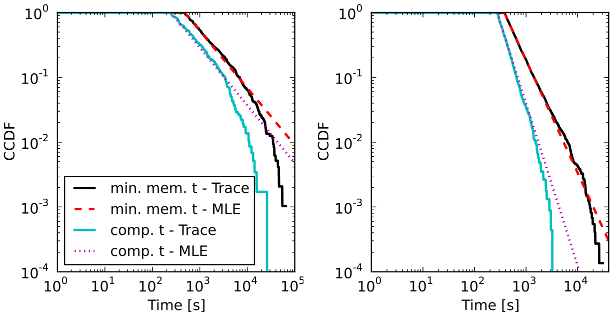

(a) Haggle trace (H06). (b) SF taxi cab trace (SF).

Figure 2: CCDF of min. membership and composition time

3. ROLE SWITCHING DISCUSSION

Given the above definition(s) of a group and its lifetime, we are now better able to tackle our original goal of fairness in the energy consumption of connected groups. For the sake of simplicity, we focus here on simple, "equal share" fairness, where each node ideally uses the same amount of energy on the long term.

As mentioned in the introduction, the challenge of any role switching scheme is to find a good trade off between frequent switches (inefficient, due to one-time switching costs) and rare switches (risking unfairness, due to limited group lifetime). Since node mobility and, hence, group lifetimes are non-deterministic, the switching interval must be re-evaluated at each switch, according to the predicted *remaining group lifetime*. As shown in [6], for a Pareto group lifetime, the remaining group lifetime CDF after t_{elapsed} is:

$$F_T(t) = 1 - \left(\frac{t_{\text{elapsed}}}{t_{\text{elapsed}} + t}\right)^{\alpha}. \tag{1}$$

We can then use, for example, the median as a prediction of a typical remaining group lifetime, which is given as:

$$t_{\text{med}} = t_{\text{elapsed}} \cdot \left(2^{\frac{1}{\alpha}} - 1\right). \tag{2}$$

As shown in Fig. 3, the median remaining group lifetime increases linearly with the groups age (elapsed contact time). Thus, the hosting role should be passed on to the next node after $\frac{t_{\text{med}}}{n}$ in a group of n nodes to achieve equal time shares.

However, a role switch incurs a non-negligible energy cost to the group of n connected nodes (1 host and $n-1$ clients) of:

$$E_{\text{switch}} = E^{HC} + E^{CH} + (n-2) \cdot E^{CC},$$

where E^{HC} is the cost to switch from host to client, E^{CH} from client to host and E^{CC} for one client to switch between two hosts. Thus, for a switch to make sense from the perspective of the whole system, the global imbalance in consumed energy among the nodes must be at least equal to E_{switch}. The global imbalance after a time t_{elapsed} is the *extra* energy the host invested in order to keep the system connected. This amounts to $t_{\text{elapsed}} \cdot (P^H - P^C)$, where P^H is the host power and P^C the client power[3].

Summarizing, the minimum time after which it makes sense to switch is $t_{\min} = \frac{E_{\text{switch}}}{P^H - P^C}$, which implies

$$t_{\text{switch}} = \begin{cases} \frac{t_{\text{med}}}{n}, & \text{if } \frac{t_{\text{med}}}{n} > t_{\min} \\ t_{\min}, & \text{else.} \end{cases} \tag{3}$$

This simple algorithm we sketched works for the intuitive "equal share" fairness. Adapting it to other definitions of fairness, such as the metric of Jain et al. [13], is part of our ongoing work.

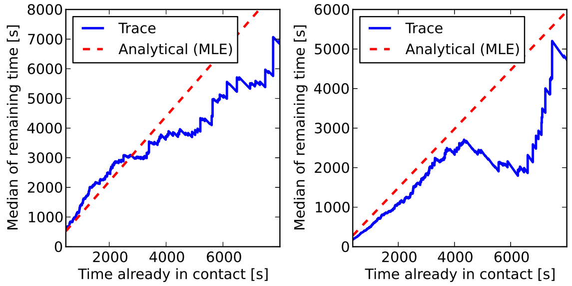

(a) Haggle trace (H06). (b) SF taxi cab trace (SF).

Figure 3: Median of remaining lifetime depending on the time already connected.

4. CONCLUSIONS AND FUTURE WORK

Energy-efficient operation is a key prerequisite for user acceptance of opportunistic device-to-device communication. One issue with widespread device-to-device communication technologies is their uneven energy consumption among connected devices, in function of their roles (host or client). This can be solved by periodically switching roles during contacts. In order to design efficient role switching schemes, in this paper, we defined and analyzed the lifetime of groups of connected devices. Based on the type of group lifetime distribution found in real-world contact traces, we sketched a simple role switching scheme aimed at achieving even energy consumption among the connected nodes of a group. In continuation of this work, we are evaluating the proposed role switching scheme, as well as investigating other potentially interesting ways of sharing energy consumption among connected nodes.

5. ACKNOWLEDGMENTS

This work is funded by ETH Zurich's Grant ETH-20 09-1 and by the EU Commission's SCAMPI (FP7-258414) FIRE Project.

6. REFERENCES

[1] Chaintreau A, Hui P et al. *Pocket switched networks: real-world mobility and its consequences for opportunistic forwarding*. Tech. rep., University of Cambridge, 2005.

[2] Lenders V, May M et al. *Wireless ad hoc podcasting*. ACM SIGMOBILE Mob Comput Commun Rev, 2008.

[3] Guo S, Falaki MH et al. *Very low-cost internet access using KioskNet*. ACM SIGCOMM Comput Commun Rev, 2007.

[4] Hossmann T, Carta P et al. *Twitter in disaster mode: security architecture*. ACM SWID. 2011.

[5] Han B, Hui P et al. *Cellular traffic offloading through opportunistic communications: a case study*. ACM CHANTS. 2010.

[6] Trifunovic S, Picu A et al. *Slicing the battery pie: Fair and efficient energy usage in device-to-device via role switching*. CHANTS. 2013.

[7] Chaintreau A, Hui P et al. *Impact of Human Mobility on Opportunistic Forwarding Algorithms*. IEEE TMC, 2007.

[8] Heimlicher S and Salamatian K. *Globs in the primordial soup: the emergence of connected crowds in mobile wireless networks*. ACM MobiHoc. 2010.

[9] Sarafijanovic N, Piorkowski M et al. *Island hopping: Efficient mobility-assisted forward. in partitioned nets*. IEEE SECON. 2006.

[10] Lenders V, Wagner J et al. *Measurements from an 802.11b mobile ad hoc network*. IEEE EXPONWIRELESS. 2006.

[11] Piorkowski M, Sarafijanovic N et al. *A parsimonious model of mobile partitioned networks with clustering*. IEEE COMSNETS. 2009.

[12] Clauset A, Shalizi C et al. *Power-law distributions in empirical data*. SIAM Review 51, 2009.

[13] Jain R, Chiu DM et al. A quantitative measure of fairness and discrimination for resource allocation in shared computer system. Eastern Research Laboratory, Digital Equipment Corporation, 1984.

---

[3] For Bluetooth, where the clients use more energy than the host [6], this formula must be adapted accordingly.