

Passing the Torch: Role Alternation for Fair Energy Usage in D2D Group Communication

Sascha Trifunovic, Andreea Hossmann-Picu, Theus Hossmann
Computer Engineering and Networks Laboratory
ETH Zurich, Switzerland
{trifunovic, picu, hossmann}@tik.ee.ethz.ch

Abstract—By using device-to-device (D2D) communication, opportunistic networks promise to fill the gaps of the networking infrastructure in remote areas, to enable communication in emergency situations, and to inspire new applications. Yet, to become feasible in practice and accepted by users, it is crucial that the energy costs of D2D connections are small and shared fairly. Fairness, in particular, is a major issue with today’s D2D technologies (Bluetooth, Wi-Fi Direct): since each connected peer must assume one of two different roles – access point/client, master/slave, the energy consumption inside a connected group is very asymmetric. While a large body of research exists on role assignment and topology control, the above issue of energy fairness is either not at all addressed (e.g. in the context of Bluetooth scatternets) or is addressed under fundamentally different conditions (e.g. in very dense and often static wireless sensor networks).

In this paper, we tackle the fairness problem of the energy consumed in a group of D2D-connected nodes, by using role switching: the two types of roles are alternated among group members, thus producing a fairer cost sharing. First, we analyze contact traces for their group topologies and find that four simple motifs – *clique*, *star*, *chain* and *NxM-clique* – cover up to 94% of the aggregated lifetime of all connected groups. We then determine the optimal role switching strategies for these motifs by formulating the cycle of role assignments as an optimization problem. Since deriving the optimal cycle online, in a distributed manner is hardly possible in practice, we also propose two role switching heuristics for online use: a randomized switching scheme tunable for efficiency or fairness, and a deterministic scheme which additionally guarantees the group’s connectivity. Finally, we evaluate our solutions on real contact traces and show that our heuristics find very good points of operation in the fairness–efficiency tradeoff.

I. INTRODUCTION

Whenever a group of users carrying smartphones move within transmission range of device-to-device (D2D) communication technologies (e.g., Bluetooth or Wi-Fi Direct), they have the opportunity to connect and exchange data without relying on any networking infrastructure. Such *opportunistic networks* [1], [2] are a promising extension to traditional networks in remote and rural areas [3], in case of broken or censored infrastructure [4], or to offload traffic from the frequently overloaded cellular network [5]. Furthermore, opportunistic networks are a more natural support for novel applications, such as highly local communication [6].

Yet, to make opportunistic networking attractive for the users of resource-constrained devices (e.g. smartphones), it is crucial to optimize its energy consumption. To this end, research efforts have mainly focused on reducing the energy used by background operations, such as neighbor discovery [7], [8]. State-of-the-art neighbor discovery adapts to the user’s context and minimizes the number of required scanning operations, while detecting a maximum number of

communication opportunities. Although these are important contributions towards *energy efficiency*, we raise a different issue here: that of *energy fairness*. Whenever a node pair or a larger group is connected, the involved devices take different roles. In Bluetooth, a device is either in master or in slave state, whereas in Wi-Fi Direct, a device can be access point (AP) or client. Measuring the energy consumption of these different roles in Bluetooth, Wi-Fi Direct and WLAN-Opp¹ [9], we find that the two roles differ by up to a factor of 5 in terms of energy consumption [10], which results in substantial unfairness when the roles are assigned statically.

To mitigate this unfairness, we propose that the nodes of a connected group periodically switch their roles. While a simple round robin scheme solves the problem for groups of trivial topology (e.g. pairs or larger *cliques*), we observe in a detailed analysis of four measured contact traces, that nodes are routinely connected in more complex topologies: *chains*, *stars*, *chained cliques*, etc (up to 68% of the aggregated lifetime of all connected groups). In such topologies, more than one master (or AP) may be required to enable all pairs of devices to communicate with each other; further, not all devices are equally well positioned for such a connection hosting role. Hence, the choice of the cycle in which devices switch roles becomes a challenging problem of finding a good balance between fairness (relieving the best located devices from always hosting the connections) and efficiency (selecting the smallest possible number of hosts).

The issue of assigning host/client roles on arbitrary topologies to form a connected group or network is well investigated in mobile ad hoc networks. More specifically, while our work is not technology specific (it is valid for any protocol with host/client roles), Bluetooth scatternet formation (BSF) [11] exhibits the most similarity to our context, in terms of the basic constraints on network formation. The explicit goal of BSF role assignment schemes varies: most aim merely at (efficient) connectivity [12], [13], some want to maximize path diversity [14], few others optimize capacity or throughput. The means to achieve these goals range from the simple, centralized solution [12], to more sophisticated algorithms based on finding (connected) dominating sets [14] or building trees [13]. However, none of these solutions address the problem of unequal energy consumption, stemming from static role assignments.

While BSF research had to work around the pre-specified roles, in wireless sensor networking (WSN), roles have been intentionally introduced to improve energy efficiency and network lifetime [15]. In WSNs, each node transmits the

¹WLAN-Opp is a method based on traditional WLAN: a subset of devices are in access point mode, such that the others can connect.

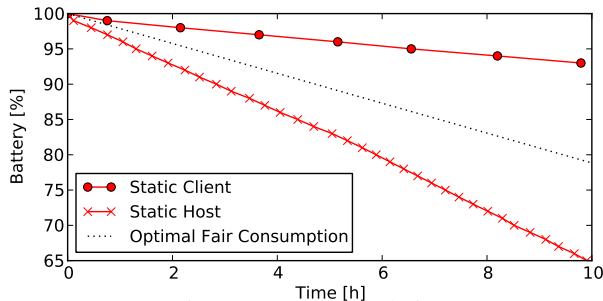


Fig. 1: Battery depletion.

same type of data to a common sink, such that aggregating the data is most efficient; to this end, node clusters are created and cluster heads elected, forming a backbone for the network. However, the energy intensive roles of cluster heads result in unfair energy usage and premature network outage. To cope with this, numerous clustering algorithms rotate the cluster role in a randomized fashion, pioneered by LEACH [16]. Another load balancing scheme, proposed by Amis and Prakash [17], works for any clustering algorithm based on dominating sets (DS). They achieve a round robin role alternation, by introducing virtual IDs that change in each round. However, both the randomized LEACH-based schemes and the ones for DS-based clustering have restrictive assumptions on the underlying topology (e.g. very dense or even fully meshed topology), which make them unsuitable for our arbitrary topologies. Finally, other popular approaches are based on alternating through various connected dominating sets (CDS) [18], [19], to provide the network with a backbone and satisfy all connectivity constraints. However, the fact that all hosts must be connected in this case is an overconstraint for our context, putting an unnecessary limit on the achievable fairness-efficiency energy tradeoff.

Our work fills the gap of these two topics by introducing three algorithms to achieve any desired fairness-efficiency tradeoff in host-client-based D2D communication on an arbitrary underlying topology. As a first step, we formulate the cycle of role assignments as an optimization problem similar to Ajmone et al. [20]. However, instead of optimizing the load on the topology, the objective function we maximize is a tunable linear combination of efficiency (percentage of nodes that have the hosting role at any given time) and fairness (here we use equality fairness [21], but the approach is easily generalizable to other types of fairness). To preserve the group’s topology regardless of its nodes’ roles, we require that at each time instance, the nodes in the hosting role form a *Neighborhood Connected Dominating Set (NCDS)*² [22]. To find the NCDSs, we determine the *edge clique cover* [23] and require that each clique has at least one host. Realizing this connection between the NCDS and edge clique covers allows us to simplify the constraints of the optimization problem to one simple condition (instead of eleven [20]). Finally, the optimal role alternation strategy is given by the cycle of NCDSs which maximizes the fairness-efficiency tradeoff.

Operating the above optimization scheme online, in a distributed way is difficult in practice as it requires global coordination. However, based on the gained insights, we propose two online role switching heuristics (one randomized and one deterministic), both tunable for efficiency or fairness. For these

²All NCDS are a subset of all DS that contain all CDS, i.e., $CDS \subseteq NCDS \subseteq DS$.

Abbr.	State/Operation	Power/Energy	STD
P_{BT}^C	Bluetooth connected (slave)	58.49 mW	3.29 mW
P_{BT}^H	Bluetooth connected (master)	28.53 mW	0.05 mW
P_O^C	WLAN-Opp associated (station)	60.79 mW	9.74 mW
P_O^H	WLAN-Opp associated (AP)	210.97 mW	11.72 mW
P_D^C	Wi-Fi Direct connected (station)	49.75 mW	3.90 mW
P_D^H	Wi-Fi Direct connected (AP)	231.92 mW	9.14 mW

TABLE I: Power consumption of connection operations.

heuristics we show that they can find near-optimal solutions for the above four classes of topologies (or *motifs*): cliques, stars, chains and $N \times M$ -cliques (N -clique connected to M -clique with x overlapping nodes), which together represent 44 – 94% of aggregate group lifetime in the contact traces. Finally, we also evaluate the heuristics’ overall performance by applying them to contact traces, which also include all kinds of more “exotic” topologies. We find that one of our heuristics can give us 89 – 98% fairness while being 88 – 97% efficient when using the most energy imbalanced Wi-Fi Direct protocol.

II. MOTIVATION

In this section, we motivate our work from two viewpoints: First, we briefly describe each D2D technology, detailing its required communication roles and their different energy consumptions. This illustrates well and quantifies the fairness problem. Second, opportunistic networking studies usually assume that nodes always connect in pairs, in which case the fairness problem is easily solved by rotating the communication roles via a simple round robin. By analyzing four real contact traces for connectivity of larger groups, we show that more complex topologies (for which the round robin solution is insufficient) are present in rather large amounts.

A. The Unequal Roles in D2D Connections

Three technologies are currently available to establish a high throughput, mid-range, D2D communication opportunity: Bluetooth, Wi-Fi Direct, and WLAN-Opp, a method based on traditional WLAN access point and station functionality. All three technologies define two communication roles: a connection *host* and a *client*.

Bluetooth intends to provide wireless connectivity in energy constrained personal area networks. In order to enable communication, one device will become the master of the connection(s) (host), serving up to seven slaves (clients), to form a piconet. Slaves may be part of multiple piconets and the master of one piconet may be a slave to another piconet’s master, thus bridging the two piconets to form a scatternet.

Wi-Fi Direct is Wi-Fi Alliance’s response to Bluetooth offering longer range and higher throughput. Two or more devices communicate by having one device in soft AP mode (host), while all others connect to it as stations (clients).

WLAN-Opp is a custom protocol, based on traditional WLAN AP and station functionality [9], which establishes communication similar to Wi-Fi Direct.

Unsurprisingly, the energy required to be a host is not equal to that needed to be a client, leading to unfair battery drain. This is the case both for the connection establishment phase, and to an even higher degree for connection maintenance and actual traffic. More precisely, the device that *hosts* the connection incurs a much higher energy cost than a *client*

	H06	MIT	ETH	SF
# contacts	128 979	75 425	22 958	1 339 274
# nodes	78	96	20	536
scanning interval	2 min	5 min	2 s	30 s
total connection t	671.07h	7921.84h	101.41h	8847.86h
clique	72.75%	72.26%	31.47%	83.37%
chain	9.77%	10.52%	3.37%	6.38%
star	2.23%	1.55%	2.00%	0.06%
NxM-clique	2.56%	4.75%	6.79%	3.71%
other	12.69%	10.92%	56.37%	6.47%

TABLE II: Motif presence (% of total connected time) in contact traces.

device³. This effect is illustrated in Fig. 1, which reports the battery usage of two phones (Samsung Galaxy Nexus with Android 4.2 in airplane mode) connected to each other via WLAN-Opp for 10 hours without transmitting data. As made obvious in this figure, the host’s battery drains much faster than the client’s – an inequitable situation.

More detailed measurements⁴ of this inequality are presented in Table I. We measure the power, P , a device continuously consumes while being a host (P^H) or a client (P^C). The device’s role impacts energy consumption for all technologies by a factor varying between 2 and 5.

To put these numbers into the perspective of the 6.48 Wh battery life of a Galaxy Nexus: The Wi-Fi Direct AP state consumes a whopping 3.58 battery-percent per hour (%/h), while the station state only consumes 0.77%/h.

B. Connected Groups in Real Traces

While the above fairness problem is easily solved for pairs of connected nodes (by switching between the two roles via round robin), finding a solution for larger groups of nodes, connected in arbitrary ways is much more challenging. In the following, we investigate whether such larger connected groups are indeed present in the four following real contact traces (also summarized in the upper half of Table II):

H06 the Hagle 2006 trace, collected during the three days of the Infocom conference in 2006 [24];

MIT the MIT reality mining trace, collected from students and staff on the MIT campus during several months [25];

ETH the ETH trace, collected from iPAQ carrying researchers at the ETH Zurich campus [26];

SF the San Francisco taxi cab trace recorded GPS position for over 500 taxicabs over a period of a month [27].

For all of the above traces, we extract the different ways in which a group of users connects most frequently. We call the most common topologies *motifs* [28]. In order to identify and classify motifs, we sort all topologies found in the traces by the total time they exist in each trace. A list of the most dominant topologies is shown in Table III. We quantify the relative importance of each topology via the fraction of the total trace time when that topology exists.

Summarizing from Table III, we identify four distinct motifs which capture the majority of the traces’ topologies. The motifs – *cliques*, *chains*, *stars*, and *NxM-cliques* – and their

³The opposite is true for Bluetooth: a master uses less energy than a slave. However, the unfairness problem remains unchanged.

⁴We took exact measurements with the Monsoon Power Monitor, which replaces the battery to record in real-time, the power consumed by the device at a resolution of 500 Hz (or 2 ms). Different device models/brands show qualitatively similar results.





















Time [% of total time]				Topology	Motif	Nodes
H06	MIT	ETH	SF			
66.07	59.71	19.97	75.14		Clique	2
5.05	9.46	6.04	6.43		Clique	3
8.46	9.15	2.33	5.43		Chain	3
1.24	2.22	3.60	0.98		Clique	4
0.89	1.87	1.41	1.16		NxM	4
0.86	1.26	1.16	0.73		NxM	4
0.96	1.11	0.72	0.79		Chain	4
1.42	1.01	0.98	0.06		Star	4
0.28	0.55	1.21	0.30		Clique	5
0.22	0.31	1.73	0.24		Unclassified	5
0.12	0.34	1.15	0.20		NxM	5
0.41	0.30	0.14	0.04		Unclassified	5
0.27	0.36	0.43	0.20		NxM	5
0.41	0.26	0.20	0.00		Star	5
0.13	0.26	0.46	0.25		NxM	5
0.22	0.19	1.01	0.19		Unclassified	5
0.31	0.20	0.25	0.13		Chain	5
0.11	0.23	0.36	0.17		Clique	6
0.22	0.22	0.42	0.04		Unclassified	5
0.14	0.17	0.17	0.07		Unclassified	5
87.79	89.18	43.74	92.55	Total %		
6.71	79.22	1.01	88.47	1% in h		

TABLE III: Top 20 Topologies in terms of time.

statistics are shown in the lower half of Table II. Most other topologies are a simple combination of these four motifs.

While the topologies of cliques, chains, and stars are well known, the NxM-cliques require an explanation. As the name suggests, they are composed of two cliques of sizes N and M with x overlapping nodes (we generally assume $N \geq M > x$). The total number of nodes in this motif is $N + M - x$.

Its interesting to note that the distribution of motifs depends on the context of the trace. The more static a trace is, the more stable complex topologies become. For instance, in the ETH trace, the contacts are strongly influenced by the office layout, resulting in unclassified topologies in over half of the time.

Considering the energy fairness problem in such complex groups is clearly more challenging than in the simple case of a connected node pair. In the next section, we rigorously define our problem as well as all the concepts needed to solve it.

III. PRELIMINARIES AND DEFINITIONS

In the following, we first illustrate the fundamental challenges we are faced with in addressing the fairness problem of D2D group communication, while still keeping overall energy consumption in check. Then, we formally define measures for fairness and efficiency in our setting, and discuss their tradeoff.

A. Constraints for an Equalizing Solution

As discussed above, an obvious remedy to the energy inequality highlighted in Section II-A would be for the connected devices to regularly switch roles and share the duty of the host. While role switching may be straightforward in the case of a node pair (the typical opportunistic *contact*), it is very challenging for other topologies, such as the ones we have just discovered in the traces.

Definition 1 (Topology). *The topology of a group of connected nodes is determined solely by their physical proximity: an edge between two group members denotes the fact that they are within transmission range of each other.*

With today's D2D technologies, operating an arbitrary, connected topology, such as the one in Fig. 2, means that each node must take on one of the two roles (host or client), forming a *configuration*.

Definition 2 (Configuration). *For a given n -node topology, a configuration or assignment of roles to nodes is represented by a binary vector $\mathbf{s} = (s_i)_{1 \leq i \leq n}$, where $s_i \in \{0, 1\}$ is the role of node i : a client if $s_i = 0$ or a host if $s_i = 1$.*

In order for a given configuration to succeed in preserving the original topology (in the sense of Def. 1), the D2D technology in use must fulfill several conditions:

- (i) multiple devices may take a hosting role;
- (ii) every client may connect to multiple hosts;
- (iii) every host may also simultaneously be a client (at no added energy cost), and thus connect to other hosts.

Under these conditions, a topology is preserved *if and only if* the set of hosting nodes (masters or APs) forms a neighborhood connected dominating set (NCDS) [22].

Definition 3 (NCDS). *A neighborhood connected dominating set is a special case of a DS, where the induced subgraph of the nodes in the set and their neighbors is connected.*⁵

Using the above defined NCDS concept, the task of role switching is reduced to finding *at least* two appropriate NCDSs for a given topology and alternating among them, so as to ensure an equitable battery depletion for all group members.

Note that, in addition to being an NCDS, the roles configuration must also be *efficient* in terms of overall energy consumption. This means that there should be as little hosts as possible, while still forming an NCDS. In the example of Fig. 2, the most efficient host configuration is formed by nodes 4 and 5, as any other NCDS requires at least three hosts. However, achieving fair energy use requires switching among several NCDSs, some of which might be less efficient. In the following, we discuss this inherent fairness–efficiency tradeoff.

B. The Fairness–Efficiency Tradeoff

When dealing with arbitrary group topologies, the choice of which nodes should be configured as a host (master or AP) confronts us with a tradeoff between efficiency and fairness. To illustrate this tradeoff, we analyze the extreme case of a star topology shown in Fig. 3. In this topology, the most efficient configuration (with respect to overall energy use) is to set the center node as the host and the rest as clients. This, however, is not particularly fair. To increase fairness, each of the outer nodes should eventually also become a host, as in the alternative NCDS. This, in turn, is not particularly efficient. Depending on the required tradeoff, the presented star topology should spend more time in the efficient one host configuration or in the fairer six host configuration.

In order to be able to rigorously quantify and tune this tradeoff, in the following, we formally define both fairness and efficiency in relation to our role switching problem.

⁵Not every DS is neighborhood connected: in Fig. 2, nodes 3 and 7 form a dominating set, but the induced subgraph of 3, 7 and their neighbors misses the edge (4, 5) to be connected. On the other hand, every *connected dominating set* is also neighborhood connected; however, the converse is not true.

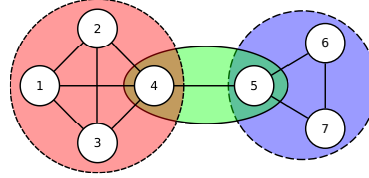


Fig. 2: Sample topology and its edge clique cover.

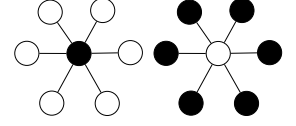


Fig. 3: Star topology with different host configurations.

Fairness Measure: Fairness, in its simplest form, corresponds to equal consumption of energy. However, in more formal settings, several measures of fairness are available, each with different goals. For example, if the group's interests are more important than those of individual devices, then devices with more energy should also contribute more, thus maximizing the lifetime of the group. Another example is the existence of special devices in the group, that is devices providing a service required by all others. In this case, the special devices should use as little energy as possible, so as to preserve the availability of the provided service.

For a generic fairness measure, applicable to any scenario, we set the goal of achieving an arbitrary division of energy usage among the connected devices. We denote this division with the vector $\mathbf{a} = (a_i)_{1 \leq i \leq n}$, for a group of n connected nodes. Each element $a_i \in (0, 1)$ denotes node i 's desired fraction out of the total amount of power, P , and the sum of all a_i 's must be 1. Further, denoting by $\mathbf{P} = (P_i)_{1 \leq i \leq n}$ the *actual* power costs incurred by each device (with $P = \sum P_i$), our fairness goal is to drive the P_i 's as close as possible to the $a_i P$'s.

While there are different ways of quantifying the perception of fairness, the most generic and intuitive metric was proposed by Jain et al. [21]. The Jain fairness index has some very important properties: it is independent of the number of devices, of the scale, and of the energy consumption measure. The resulting fairness value is bounded between 0 and 1 and continuous, ensuring that it varies discernibly with every change of the ratio of the consumed power.

The Jain fairness index $\mathcal{J}(\mathbf{a}, \mathbf{P})$ for power consumption among n devices is defined as follows:

$$\mathcal{J}(\mathbf{a}, \mathbf{P}) = \frac{\left(\sum_{i=1}^n \frac{P_i}{a_i} \right)^2}{n \cdot \sum_{i=1}^n \left(\frac{P_i}{a_i} \right)^2}. \quad (1)$$

With the above definition of an allocation metric, the index measures the deviation of the vector of consumed powers $\mathbf{P} = (P_i)_{1 \leq i \leq n}$ from the desired division $\mathbf{a}P = (a_i)_{1 \leq i \leq n}P$. The denominator of Eq. (1) is minimal (and thus equal to the nominator), whenever all the summands are equal. The denominator increases with the summands' variations, reducing the overall fairness.

Efficiency Measure: Conflicting with a fairness objective is usually an efficiency goal. In our case, some host configurations comprise less hosts than others, thus consuming less energy overall and being more efficient.

To be able to set a desired tradeoff we need to quantify efficiency. We measure the efficiency of a configuration as the ratio of how close this configuration is to the optimally efficient one. More specifically, for every group of n nodes, there is a minimal number of hosts h_{\min} that are required to achieve the connected underlying topology, in the sense of

Def. 1. This host configuration is the minimal NCDS of that topology, and the power it consumes is given by:

$$P_{\min} = h_{\min} \cdot P^H + (n - h_{\min}) \cdot P^C. \quad (2)$$

Then, the efficiency of an arbitrary configuration \mathbf{s} , with $h_{\mathbf{s}} = \sum_{i=1}^n s_i$ hosts can be described as:

$$e(\mathbf{s}) = \frac{P_{\min}}{h_{\mathbf{s}} \cdot P^H + (n - h_{\mathbf{s}}) \cdot P^C}. \quad (3)$$

The efficiency $e(\mathbf{s})$ is a value in $(0, 1]$, where 1 means perfect efficiency and 0 the opposite.

Concluding this section, we now have a clear picture of the problem and are equipped with formal definitions of the tools needed to solve it.

IV. ROLE SWITCHING ALGORITHMS

Considering the above constraints, as well as our fairness and efficiency measures, in this section, we propose three algorithms for role switching in arbitrary node topologies: a centralized optimal scheme and two distributed heuristics.

Let \mathcal{N} be our group of n connected nodes, forming some arbitrary topology, such as the one in Fig. 2. At any given time t , the topology is operated (via one of the D2D technologies) with a role configuration $\mathbf{s}(t) = (s_i)_{1 \leq i \leq n}$, where a node i is a host for $s_i = 1$ or a client for $s_i = 0$. To achieve a certain fairness objective, we need a set \mathcal{M} of distinct configurations to cycle through, with $|\mathcal{M}| = m \geq 2$. According to the constraint discussion in the previous section, each of these configurations must be an NCDS.

To summarize, assuming slotted time t and given:

- (i) the set of all NCDS role configurations for \mathcal{N} , and
- (ii) a fair energy allocation vector $\mathbf{a} = (a_i)_{1 \leq i \leq n}$,

our challenge is to find the cycle \mathcal{M} of role configurations, which achieves the desired fair allocation without disregarding efficiency. To operate the topology, the group of n nodes can then, at each timeslot, cycle through the configurations in \mathcal{M} .

We note that, the problem of finding NCDSs is still new and not so well investigated. Studying solutions to this problem is beyond the scope of our work; here, to find a topology's NCDSs, we use a simple scheme whereby we first identify the largest possible cliques that together cover all the edges of that topology, i.e. an *edge clique cover* [23]. Then, NCDSs can be built by ensuring that each clique contains at least one NCDS member. Applying this to Fig. 2, we obtain three cliques: $\{1, 2, 3, 4\}$, $\{4, 5\}$, $\{5, 6, 7\}$, and the nodes 4 and 5 form a minimal NCDS.

This connection between the edge clique cover and the NCDS has two important advantages: (i) In a *centralized case*, we can extract all maximal cliques with the Bron-Kerbosch algorithm [29] and then formulate the configuration constraints with one simple equation. (ii) In a *distributed case*, the nodes do not have to know the whole topology but only the cliques they belong to. This can be easily calculated by each node knowing all the 2-hop neighbors [30].

A. Optimal Role Switching (OPT)

To find the optimal solution \mathcal{M} for the above role switching problem, we can formulate the desired fair energy allocation and the efficiency into an optimization problem. To reconcile the conflicting goal of maximizing both fairness and efficiency, we aim at maximizing the following objective function, which

corresponds to exactly one cycle through \mathcal{M} , requiring m timeslots:

$$f_{\text{obj}}(\mathcal{M}) = e(\mathcal{M}) + \gamma \cdot \mathcal{J} \left(\mathbf{a}, \left(\sum_{\mathbf{s} \in \mathcal{M}} P_i(\mathbf{s}) \right)_{1 \leq i \leq n} \right), \quad (4)$$

where the γ term tunes the fairness–efficiency tradeoff, and $P_i(\mathbf{s}) = s_i P^H + (1 - s_i) P^C$ is the power consumed by node i in role configuration \mathbf{s} . Assuming equality fairness, for simplicity of illustration, and substituting the definitions of fairness and efficiency⁶, the objective function becomes:

$$\begin{aligned} f_{\text{obj}}(\mathcal{M}) &= \frac{m P_{\min}}{\sum_{\mathbf{s} \in \mathcal{M}} h_{\mathbf{s}} P^H + (n - h_{\mathbf{s}}) P^C} + \gamma \cdot \frac{\left(\sum_{i=1}^n \sum_{\mathbf{s} \in \mathcal{M}} P_i(\mathbf{s}) \right)^2}{n \cdot \sum_{i=1}^n \left(\sum_{\mathbf{s} \in \mathcal{M}} P_i(\mathbf{s}) \right)^2} \\ &= \frac{m P_{\min}}{\sum_{i=1}^n \sum_{\mathbf{s} \in \mathcal{M}} P_i(\mathbf{s})} + \gamma \cdot \frac{\left(\sum_{i=1}^n \sum_{\mathbf{s} \in \mathcal{M}} P_i(\mathbf{s}) \right)^2}{n \cdot \sum_{i=1}^n \left(\sum_{\mathbf{s} \in \mathcal{M}} P_i(\mathbf{s}) \right)^2} \end{aligned} \quad (5)$$

To complete our optimization problem, we add the constraint that all elements of \mathcal{M} should also be NCDSs. Given our above method of finding NCDSs based on an edge clique cover, this means that for all $\mathbf{s} \in \mathcal{M}$ the node states s_i must be such that there is at least one host per clique in the given topology. Let $K = \{K_1, K_2, \dots, K_l\}$ be the topology's set of maximal cliques⁷. Then our optimization problem is:

$$\begin{aligned} \max_{\mathcal{M}} f_{\text{obj}}(\mathcal{M}) \text{ such that} \\ \sum_{i \in K_k} s_i \geq 1, \quad \forall K_k \in K \text{ and } \forall \mathbf{s} \in \mathcal{M}. \end{aligned} \quad (6)$$

This is a mixed integer optimization problem with a non-convex objective function. As there is no out-of-the-box software to solve such problems, we transform our objective function to make it convex, while maintaining the same solution space and the same global maxima. Maximizing the efficiency term defined in Eq. (3) is equivalent to minimizing its inversion (ignoring the constant terms). For the fairness term, we proceed as follows. Fairness is maximized when all nodes use the same amount of energy. We thus minimize the the sum of all energy differences among all possible combinations of nodes. This results in the following convex objective function to be *minimized*:

$$g_{\text{obj}}(\mathcal{M}) = \sum_{i=1}^n \left(\sum_{\mathbf{s} \in \mathcal{M}} P_i(\mathbf{s}) + \gamma' \cdot \sum_{j=i+1}^n \left| \sum_{\mathbf{s} \in \mathcal{M}} (P_i(\mathbf{s}) - P_j(\mathbf{s})) \right| \right). \quad (7)$$

Finally, minimizing $g_{\text{obj}}(\mathcal{M})$, we obtain an optimal set of role configurations \mathcal{M}_{opt} . When our group of nodes \mathcal{N} cycle through these configurations, they achieve the requested tradeoff between fairness and efficiency, while preserving a connected topology.

In addition to finding practically feasible cycles \mathcal{M} , it is also of interest for evaluation purposes to calculate the theoretically

⁶Note that the efficiency of a set of role configurations is not the sum of individual efficiencies.

⁷We use the Bron–Kerbosch algorithm [29] to find these.

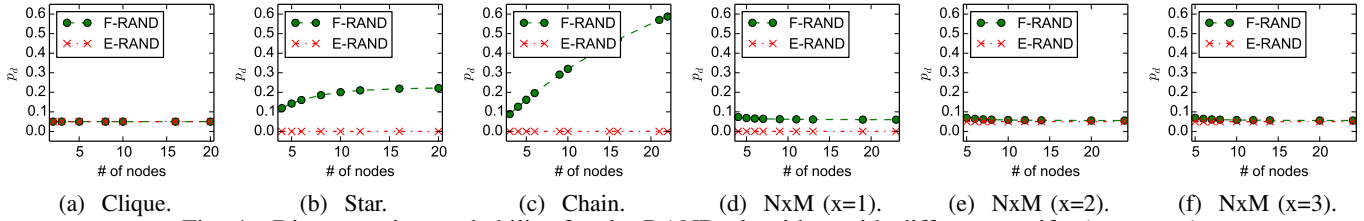


Fig. 4: Disconnection probability for the RAND algorithm with different motifs. ($p_d = 0.05$)

optimal fairness–efficiency tradeoff. This can be done by relaxing our binary state variables $\mathbf{s} = (s_i)_{1 \leq i \leq n}$ via convex relaxation [31]. Solving the optimization problem under relaxation results in a single (unfeasible) role configuration vector \mathbf{s}_{opt} , with elements in $(0, 1)$ rather than $\{0, 1\}$. Then, $f_{\text{obj}}(\mathbf{s}_{\text{opt}})$ is the desired optimal theoretical tradeoff. This allows us to determine the whole space of possible tradeoffs as a reference for other algorithms.

Summarizing, this section has introduced a scheme for finding the optimal cycle of role configurations \mathcal{M} for a given topology \mathcal{N} , given a desired tradeoff between fairness and efficiency. Obviously, this scheme requires that all n nodes have global knowledge about the topology which requires $O(n^2)$ messages to be exchanged. It also provides a global cycle \mathcal{M} , for globally coordinated timeslots. The distribution of the cycle requires $O(n)$ messages to be exchanged within the topology. Coordinating globally and propagating global information is extremely challenging in opportunistic networks, making this solution impractical. Therefore, we also propose two practical heuristics for the same problem in the following. We then use the optimal scheme as a benchmark for the evaluation of these heuristics in Section V-A.

B. Distributed Randomized Switching (RAND)

One simple way of avoiding the complications of global coordination and global topology knowledge is for each node to probabilistically decide, at each timeslot, whether it will be a host for that timeslot. Naturally, the hosting probability should depend on the goal of the role switching algorithm. A fair algorithm selects the same hosting probability for all nodes, while an efficient algorithm gives higher hosting probabilities to nodes that connect several cliques, as they are more likely to be part of a minimal NCDS.

Depending on the topology and the chosen hosting probabilities, there is a chance that the host nodes do not form an NCDS at a given time, leading to a disconnected topology. Determining the disconnection probability for an arbitrary topology is not an easy task. However it is more or less easily achievable for the most frequently occurring motifs from Table III. For example, in a clique of size n , all nodes are equal and should thus have equal hosting probability p_h . Then, the probability that no node is a host is given by:

$$p_d = (1 - p_h)^n. \quad (8)$$

To achieve a *desired disconnection probability* p_d , we simply invert the above equation and find the hosting probability p_h for each node in the clique:

$$p_h = 1 - p_d^{\frac{1}{n}}. \quad (9)$$

Similar, albeit more complex formulas can be obtained for the other three common types of topologies from Table III, as shown in [32]. Since these formulas are not easily invertible as

above, numerical inversion must be performed. This results in an algorithm that is opaque and difficult to tune and to evaluate beyond the aspect of disconnection probability. For a more transparent solution, we use the insight from above to devise the following two heuristics (one aimed at fairness, the other at efficiency). Both are based on local topology knowledge, that is knowledge about all two-hop neighbors.

F-RAND: In the fair variant, each node selects p_h according to Eq. 9 in its local clique. If a node is part of multiple cliques, it takes the biggest value of p_h in order not to increase the disconnection probability. This results in an overall disconnection probability that is at least equal to the desired p_d , because chaining or otherwise linking cliques increases the likelihood of disconnection as can be seen in Fig. 4 for $p_d = 0.05$. This is especially true for the chain motif, where a disconnection becomes more probable the longer the chain is. However, chains of more than five nodes are rare.

E-RAND: In the efficient variant, only nodes that bridge multiple cliques should be able to become a host, thus all non-bridging nodes set $p_h = 0$. Bridging nodes behave as follows:

- If a set of bridges connect the same set of cliques, their probability will be determined by Eq. 9 with n being the number of bridging nodes.
- If multiple nodes connect multiple cliques: bridges that serve cliques by themselves set $p_h = 1$, bridges that share cliques, see above.

Selecting p_h with such an efficiency tuned algorithm will give us overall lower disconnection probabilities than the desired p_d , especially if we have single nodes that connect multiple cliques with $p_h = 1$. This can be clearly seen in Fig. 4.

Similar to the OPT algorithm, we can tune the fairness–efficiency ratio of the RAND algorithm. If p_F is the probability a node selects with the fair variant and p_E the one selected by the efficient version, a node may become host with

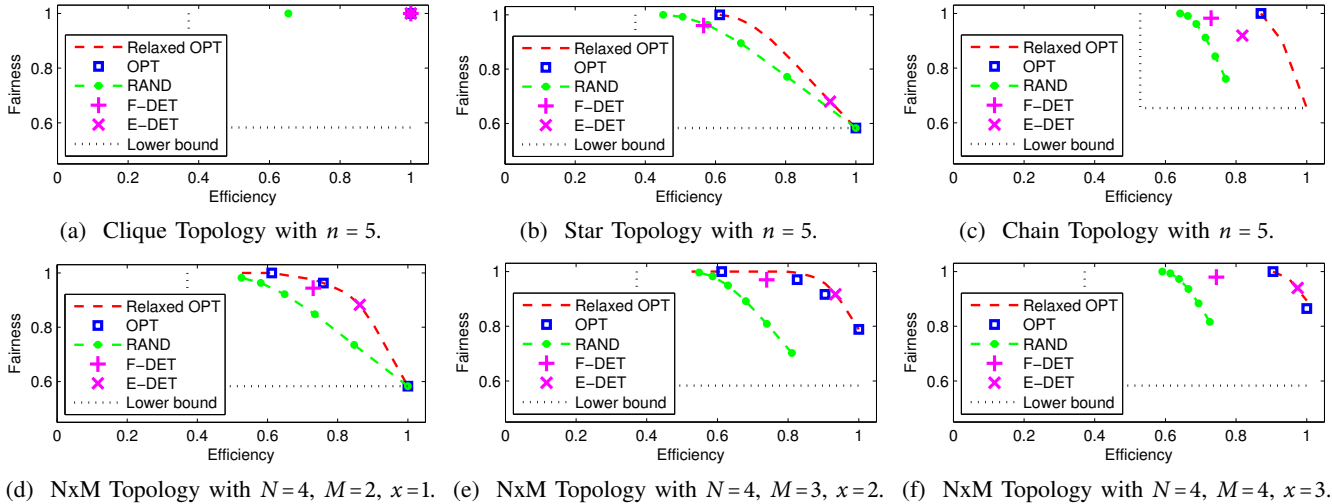
$$p_h = \beta \cdot p_F + (1 - \beta) \cdot p_E \quad (10)$$

and thus tune the ratio with parameter β .

To summarize, the random heuristic only requires local information about the 2 hop neighbors which can be exchanged with one local broadcast per node. Further, RAND does not require the creation and negotiation of a switching schedule, but it does come at the risk of a disconnected topology if the randomly chosen hosts do not form a NCDS.

C. Distributed Intra-Clique Switching (DET)

While the random host configuration is simple and does not require global knowledge or coordination among nodes, the possibility for disconnections is undesirable or might be unacceptable. For this reason, we introduce a distributed algorithm that guarantees connectedness, while only requiring local knowledge and little coordination, i.e. the broadcast of only two messages per node, one with the neighborhood information, and one with the local schedule.



(d) NxM Topology with $N=4$, $M=2$, $x=1$. (e) NxM Topology with $N=4$, $M=3$, $x=2$. (f) NxM Topology with $N=4$, $M=4$, $x=3$.
 Fig. 5: Fairness – Efficiency tradeoff for different topologies for the Wi-Fi Direct energy consumption. Note that the RAND algorithm may be more efficient than the optimum as it allows for disconnections.

In this algorithm, every clique selects a host independently of all other cliques and the hosting role is then rotated inside the clique. This eliminates the need for globally defined time-slots. However, it comes at a fairness cost: nodes belonging to multiple cliques might be scheduled to become host by two or more cliques at different times. In addition, inefficiencies may also occur: for instance, whenever a node belonging to multiple cliques is host, there might be a second unnecessary host in one of its cliques. These inefficiencies can be partly avoided by allowing neighboring cliques to share their hosting schedules and not schedule unnecessary hosts. This DET variant tuned for efficiency (E-DET), comes at the cost of fairness. In the evaluation, we show that such sporadic fairness and efficiency reductions are a small price to pay for the simplicity of this algorithm.

V. FAIRNESS–EFFICIENCY TRADEOFF EVALUATION

We now evaluate the performance of the above described algorithms, OPT, RAND, and DET, first on the four commonly observed motifs, then in the traces. For this analysis, we have implemented OPT by formalizing the optimization problem described in Section IV-A with yalmip [33] and solving it with Gurobi⁸. The number of configurations we allow in the set \mathcal{M} , is the least common multiplier of the sizes of all the cliques involved. For example, in an NxM-clique with $N = 4$ and $M = 3$, we select 12 configurations. For the RAND algorithm the chosen probabilities p_h are based on a disconnection probability of $p_d = 0.05$.

In order to determine the fairness and efficiency of each algorithm we need to calculate the fraction of time each node is expected to be in host mode which is basically defined by the probability (or fraction) to be a host p_h . For the OPT algorithm this is derived from the optimal schedule, for the RAND algorithm it is given by Eq. (10), and for the DET algorithm its expected value is calculated from the algorithms specification. Knowing the probability to find a node in the host role leads to the following expected power consumption:

$$P = P^H \cdot p_h + P^C (1 - p_h).$$

Given the power each node consumes, we can compute the fairness and efficiency from Eq. (1) and Eq. (3) respectively.

For this evaluation we use the Wi-Fi Direct energy values shown in Table I. The outcome is similar for the energy consumption of the other technologies.

A. Per Motif Evaluation

The possible fairness–efficiency tradeoff space can be seen in Fig. 5 for six topologies with $n = 5$ nodes.⁹ The marked lower bounds (dotted line) of fairness and efficiency are given by the constraints of the problem (e.g., both host and client require energy). The upper bound (red dashed line) is determined by convex relaxation of the integer optimization problem. As expected, the integer solutions of OPT are usually all located on or close to the upper bound (blue squares).

The RAND algorithm can be tuned along the green dashed line by choosing the parameter β ($0 \leq \beta \leq 1$). While in the clique (Fig. 5a), the tradeoff cannot be tweaked, the range of tradeoffs is especially large for topologies with cliques that have a single overlapping nodes, such as the star (Fig. 5b), or the NxM-clique with $x = 1$ (Fig. 5d). It is less efficient if there are multiple equally important nodes, such as in the chain (Fig. 5c), or the NxM-clique with $x = 2$ (Fig. 5e) or $x = 3$ (Fig. 5f). This is because the RAND algorithm is unable to efficiently loadbalance among equally important nodes within a topology as it has to overprovision to avoid disconnections.

The figure confirms that by introducing a tiny bit of coordination the DET algorithm generally outperforms RAND. Solutions found by DET are in almost all cases (with the exception of the chain, where efficiency is hard to achieve without global coordination) close to the optimal trade-off boundaries, which shows that the simple deterministic heuristics are effective.

B. Performance on Traces

Now we know how our algorithms perform for the most common motifs and how they compare to the optimally achievable tradeoff. However, to understand their performance in practice, we must take into account realistic traces that also contain other, more complex topologies. As explained in Section IV-B the disconnection probability also depend on the topologies, usually being higher for F-RAND and lower

⁹Slightly varying n does not change these results much. However, a more detailed analysis on how the achievable fairness–efficiency tradeoff changes with the motif size can be found in [32].

⁸www.gurobi.com

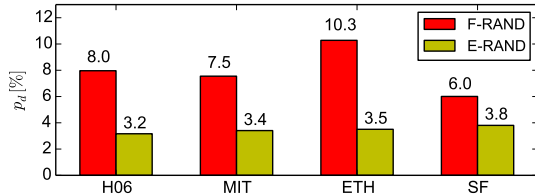


Fig. 6: Disconnection probability for RAND in traces.

for E-RAND. We can see in Fig. 6 this is the case for the actual disconnection probabilities in the four traces. The ETH trace has the most complex topologies resulting in the highest disconnection probability for the F-RAND algorithm.

The resulting tradeoffs for all traces are shown in Fig. 7. While all algorithms are generally quite fair (0.75 – 1.0), the efficiency of the RAND algorithms is not very good for all traces (0.63 – 0.86). This is because it needs to overprovision to avoid disconnections. At the cost of some local coordination, the DET algorithm (especially the efficient variant E-DET) clearly outperforms RAND without suffering from disconnections, achieving a fairness around 0.91 – 0.99 while maintaining an efficiency of 0.81 – 0.98. While DET is slightly more complex to implement as it requires some schedule coordination among local nodes, it comes close to the OPT algorithm in all traces without the need of global coordination and a costly optimization.

VI. CONCLUSION

A fair distribution of energy usage is a key prerequisite for user acceptance of opportunistic D2D communication. Current ad-hoc wireless communication technologies available for recent mobile phones, i.e., Bluetooth, Wi-Fi Direct, and WLAN-Opp, require devices to assume different roles, i.e. host or client, which differ in energy consumption by a factor of two to five. To improve fairness of the consumed energy, we propose to regularly switch roles among the members of the communicating group. However, depending on the topology of the group arbitrary switching might be inefficient or not possible. We analyze the group topologies of four contact traces and extract four motifs cover 44 – 94% of the lifetime of all connected groups. We then introduce two distributed role switching heuristics and compare them to the optimal switching schedule for the four motifs which is determined by convex optimization with global knowledge. Finally, we show that our distributed heuristics can give us 89 – 98% fairness in real world traces while achieving 88 – 97% of the maximal efficiency when using the most energy imbalanced Wi-Fi Direct protocol.

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, Karlsson G et al. *Wireless Ad Hoc Podcasting*. MC2R, 12(1):65–67, 2008.
- [3] Guo S, Falaki MH et al. *Very low-cost internet access using KioskNet*. ACM SIGCOMM CCR, 2007.
- [4] Hossmann T, Carta P et al. *Twitter in disaster mode: security architecture*. SWID, 2011.
- [5] Han B, Hui P et al. *Cellular traffic offloading through opportunistic communications: a case study*. Chants, 31–38, 2010.
- [6] Pietiläinen AK, Oliver E et al. *MobiClique: middleware for mobile social networking*. WOSN, 49–54, 2009.
- [7] Wang W, Srinivasan V et al. *Adaptive contact probing mechanisms for delay tolerant applications*. MobiCom, 2007.

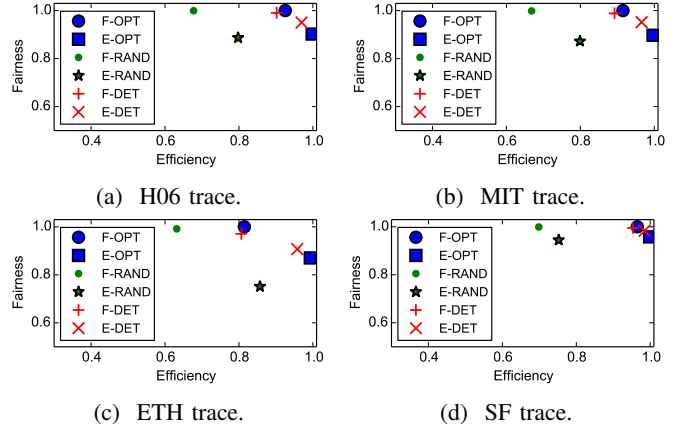


Fig. 7: Fairness – Efficiency tradeoffs (traces).

- [8] Wang Y, Krishnamachari B et al. *Markov-optimal sensing policy for user state estimation in mobile devices*. IPSN, 2010.
- [9] Trifunovic S, Kurant M et al. *WLAN-Opp: Ad-hoc-less opportunistic networking on smartphones*. Ad Hoc Networks, 2014.
- [10] Trifunovic S, Picu A et al. *Adaptive role switching for fair and efficient battery usage in device-to-device communication*. MCCR, 18(1), 2014.
- [11] Stojmenovic I and Zaguia N. *Bluetooth scatternet formation in ad hoc wireless networks*. Performance Modeling and Analysis of Bluetooth Networks: Network Formation, Polling, Scheduling, and Traffic Control, chap. 9, 2006.
- [12] Salonidis T, Bhagwat P et al. *Distributed topology construction of Bluetooth personal area networks*. INFOCOM, 2001.
- [13] Zaruba G, Basagni S et al. *Bluetrees-scatternet formation to enable Bluetooth-based ad hoc networks*. IEEE ICC, 2001.
- [14] Petrioli C, Basagni S et al. *Configuring bluestars: multihop scatternet formation for bluetooth networks*. IEEE Transactions on Computers, 52(6), 2003.
- [15] Aziz A, Sekercioglu Y et al. *A Survey on Distributed Topology Control Techniques for Extending the Lifetime of Battery Powered Wireless Sensor Networks*. Communications Surveys & Tutorials, 15(1), 2013.
- [16] Heinzelman W. *Energy-efficient communication protocol for wireless microsensor networks*. HICSS, 2000.
- [17] Amis A and Prakash R. *Load-balancing clusters in wireless ad hoc networks*. ASSET, 2000.
- [18] Chen B, Jamieson K et al. *Span: An energy-efficient coordination algorithm for topology maintenance in ad hoc wireless networks*. Wireless networks, 2002.
- [19] Bao L and Garcia-Luna-Aceves JJ. *Topology management in ad hoc networks*. MobiHoc, 2003.
- [20] Ajmone Marsan M, Chiasserini CF et al. *Optimizing the topology of Bluetooth wireless personal area networks*. INFOCOM, 2002.
- [21] Jain R, Chiu D et al. *A quantitative measure of fairness and discrimination for resource allocation in shared computer system*. Tech. rep., DEC, 1984.
- [22] Arumugam S and Sivagnanam C. *Neighbourhood total domination in graphs*. Opuscula Mathematica, 31(4), 2011.
- [23] Roberts F. *Applications of edge coverings by cliques*. Discrete applied mathematics, 10:93–109, 1985.
- [24] Chaintreau A, Hui P et al. *Impact of Human Mobility on Opportunistic Forwarding Algorithms*. IEEE TMC, 6, 2007.
- [25] Eagle N and Pentland AS. *Reality mining: sensing complex social systems*. Personal and Ubiquitous Computing, 10(4):255–268, 2006.
- [26] Lenders V, Wagner J et al. *Measurements from an 802.11b Mobile Ad Hoc Network*. IEEE EXPONWIRELESS, 2006.
- [27] Piorkowski M, Sarafijanovic N et al. *A parsimonious model of mobile partitioned networks with clustering*. IEEE COMSNETS, 2009.
- [28] Milo R, Shen-Orr S et al. *Network motifs: simple building blocks of complex networks*. Science, 298(5594), 2002.
- [29] Bron C and Kerbosch J. *Algorithm 457: finding all cliques of an undirected graph*. Communications of the ACM, 16(9), 1973.
- [30] Krishna P, Vaidya NH et al. *A cluster-based approach for routing in dynamic networks*. ACM SIGCOMM CCR, 27(2), 1997.
- [31] Boyd S and Vandenberghe L. *Convex Optimization*. Cambridge University Press, 2004. ISBN 9780521833783.
- [32] Trifunovic S, Hossmann-Picu A et al. *How to Achieve Your Fairness-Efficiency Tradeoff in D2D Communication via Role Switching*. Tech. rep., ETH Zurich, 2014.
- [33] Löfberg J. *YALMIP : A Toolbox for Modeling and Optimization in MATLAB*. CACSD, 2004.