Watts2Share: Energy-Aware Traffic Consolidation

Ekhiotz Jon Vergara, Simin Nadjm-Tehrani
Department of Computer and Information Science
Linköping University, Sweden
{ekhiotz.vergara, simin.nadjm-tehrani}@liu.se

Abstract—**Energy consumption is becoming the Achilles' heel of the mobile user quality of experience partly due to undisciplined use of the cellular (3G) transmissions by applications. The operator infrastructure is typically configured for peak performance, whereas during periods of underutilisation the handsets pay the price by staying in high energy states even if each application only uses a fraction of the maximum available bandwidth.**

In this paper we promote a bi-radio scenario where instead of independently using own cellular connections, several users share a single cellular link offered by one member of a coalition (a rotating aggregator). We present Watts2Share, an architecture for energy-aware traffic consolidation whereby group members' data flows transmitted through a second radio (e.g., WiFi) are aggregated by the aggregator and retransmitted through the cellular link. Through careful and repeatable studies we demonstrate that this scheme saves up to 68% of the total transmission energy in handsets compared to a pure 3G scenario. The studies are based on a wide range of real traffic traces and real cellular operator settings, and further illustrate that this scheme reduces the overall energy by reducing the signalling overhead, as well as extending the lifetime of all handsets.

Index Terms—**traffic consolidation; bandwidth sharing; 3G; WiFi; SoftAP;**

I. Introduction

With the advent of computationally powerful handsets we are finally entering the era of mobile services anytime anywhere. True mobility will increasingly be supported by widespread deployments of the cellular and vehicular infrastructures, as opposed to the fixed infrastructure provided by the WiFi access points.

The growing wave of mobile data communication has two extreme consequences: (1) unforeseen data volumes make the mobile operators eager to squeeze every bit of possible capacity out of their infrastructure, (2) a mass diversity of applications and platforms create an ecosystem that is hard to optimise from every conceivable perspective.

A user that has access to data/services anytime anywhere will focus on interactions with the application, expecting to communicate whenever needed without caring about battery lifetime. However, the consequence of the cellular operator's capacity optimisation, caring for maximum loads, becomes a high tax on the handsets' transmission energy during under-utilisation. More specifically, the user is left in an undesirable situation: the handset's hardware vendor will optimise the use of their energy much tighter than earlier, but their battery is wasted by the undisciplined use of data transmissions in the cellular network, using apps that were not specifically

optimised for a given radio interface, or operator setting [1], [2].

In this paper we consider the energy saving potential for users during periods in which the high energy (cellular) network is underutilised, and consider a hybrid solution that uses access to a low energy radio interface for local communication. While hybrid solutions have earlier been proposed for increasing the quality of service or extending the reach of a cellular network during *overloads* [3]–[5], we believe that a hybrid scheme to minimise the energy footprint for multiple users during *underloads* is a novel approach.

A typical 3rd generation cellular (3G) bandwidth is not fixed. Neither is the claim of a user on that bandwidth over time. However, large data traffic studies show that applications used by a single user do not always fully utilise the allocated channel bandwidth [6]. We consider the case where a user group can use a common scheme for traffic consolidation, and show that during periods of underutilisation they all save energy compared to each user connecting to a 3G network independently. The common scheme is realised by an architecture for Watts2Share: nodes form coalitions, an aggregator collects the other nodes' sporadic traffic using a low energy radio interface, and then the aggregator relays the coalition traffic via the high energy interface. The gains in energy build on maximally using the data transmission states while being connected to a 3G network.

The contributions of this paper are as follows. (1) We propose a hybrid architecture for a traffic consolidation scheme called Watts2Share, (2) exploit a tool needed to efficiently measure and quantify the footprint for a data traffic trace in this scheme (the EnergyBox), and (3) running experimental studies on a range of applications on normal user smartphones, illustrate that the cooperative scheme pays off.

The paper starts by clarifying the underlying reasons for the 3G communication energy overhead even during underloads in section II, then introduces the WiFi communication energy. After the background, the Watts2Share scheme is presented in section III and evaluated in section IV. Section V is devoted to review of related works, and section VI concludes the paper.

II. Background

This section provides the background on the mobile handset energy consumption due to packet transmission. We review the typical energy consumption when transmitting on a 3G cellular interface, and then go on to describe the handset states

when transmitting with WiFi. The measurement methodology adopted is similar to earlier works (3G [7] and WiFi [8]).

A. Energy consumption of 3G

The energy consumption of the user equipment (UE) in 3G is mostly influenced by the Radio Resource Control (RRC) and the Radio Link Control (RLC) protocols. These protocols control the radio resource allocation by the Radio Network Controller (RNC) at the network operator side. According to RRC the UE can be in different states which vary in power consumption and performance in terms of response time and data rate: Dedicated Channel (DCH), Forward Access Channel (FACH), and Paging Channel (PCH), in decreasing order.

State transition	Uplink size (bytes)	Downlink size (bytes)
PCH-DCH	850 - 1000	515
FACH-DCH	294	515
PCH-FACH	Always triggered	
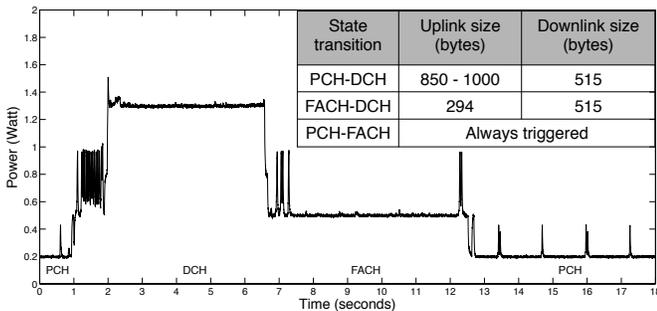
Fig. 1: Power profile of 3G and RLC buffer thresholds.

Fig. 1 shows an example plot for power consumption levels of the states in a handset when connected to TeliaSonera operator network when sending a burst of UDP packets. When the UE is in PCH, it can be paged with the lowest energy drain, but no data can be sent since the UE has no channel allocated. Before sending the UDP packets, signalling is performed to establish a dedicated channel connection and move from PCH to DCH (1-2s in Fig. 1).

The state transitions to higher performance states are triggered observing the RLC data buffers as an estimate of traffic volume. The UE reports to the RNC the observed traffic volume and if the buffer occupancy is more than a fixed threshold, the transition is triggered. The table in Fig. 1 shows the measured values for the uplink and downlink RLC data buffer thresholds [7].

Inactivity timers are used to trigger transitions to lower performance states. When the UE is in the DCH state for a duration of T_1 with small or no data transmission, the RNC releases the dedicated channel and switches the UE to FACH (7s in Fig. 1). Another timer (T_2) controls the transitions from FACH back to PCH. Inactivity timers cause an energy overhead known as energy *tail* due to the UE remaining in a high energy state while not transmitting anything [9]. This makes the energy consumption dependent on the transmission data pattern, where a single packet can trigger a move to DCH and consume both energy tails.

Fast dormancy is a mechanism that allows the UE to signal the RNC the desire to switch to a low power state before the inactivity timeout. Some networks implement a low activity

mechanism in DCH to release the transport channel and move to FACH when there is low traffic [10], and some do not.

B. Energy consumption of WiFi

In contrast to the 3G case, the WiFi station (the client device) is in control of the different WiFi states. The station is in the Constantly Awake Mode (CAM) when it has the power-saving features disabled experiencing the best performance.

The IEEE 802.11 standard defines a Power Save Mode (PSM), which allows the stations to switch to low power mode during predefined periods of time while not transferring any data. The access point (AP) buffers frames for the sleeping clients and uses the Traffic Indication Map (TIM) bitmap to indicate the presence of downlink frames at every beacon interval (typically 100 ms) to the clients. The client wakes up periodically (using multiples of the beacon interval) and sends a Power Save Poll (PS-Poll) message to the AP to proceed with the reception of each buffered frame. Fig. 2 (left) shows the PSM implementation in two smartphones and their wake up interval.
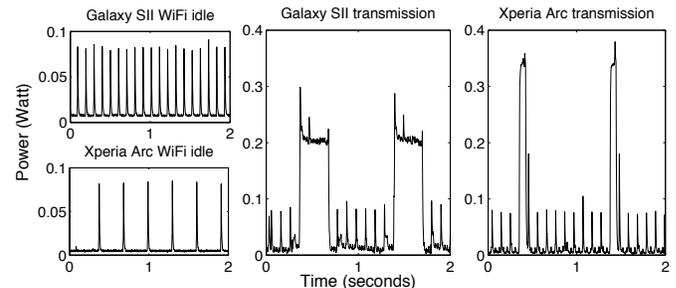
Fig. 2: WiFi energy consumption in different handsets.

Recent smartphone models implement a mechanism named Adaptive PSM to overcome the overhead and latency drawback for using the PS-Poll mechanism [11]. The client switches between the CAM and PSM modes based on heuristics (e.g., traffic inactivity period or screen on/off). Fig. 2 (center and right) shows that after a predefined inactivity timeout (δ) without packet transmission, the station switches back from CAM to PSM. This δ timeout creates an energy tail in a similar way as in 3G.

Related studies [12], [13] show that previous generation devices implement longer δ timeouts (e.g., 1.5 seconds for HTC Magic). Moreover, some drivers also implement two packets per second thresholds (Up and $Down$) that trigger the PSM-CAM and CAM-PSM transitions respectively [14].

To sum up, adaptive PSM induces a large impact on energy consumption of a WiFi station depending on the data pattern.

III. Watts2Share

A Watts2Share scenario is composed of several nodes with the goal of reducing their cellular energy consumption making use of heterogeneous networks. The actors are nodes equipped with two radio interfaces, namely cellular and short-range, that create 1-hop (star) coalitions using the short-range interface.

While reducing energy, we aspire to not reduce the users' quality of experience. Given n 3G users where each user i requests a part of the available maximum 3G bandwidth at time t ($B_i(t)$ and $B_{max}(t)$ respectively), we assume that the cumulative bandwidth claim stays under a fraction ϕ of maximum bandwidth during the regular Watts2Share operation time, where ϕ is a configurable parameter of the system:

$$\sum_{i=1}^{n} B_i(t) < B_{max}(t) \cdot \phi \tag{1}$$

Traffic consolidation in Watts2Share is designed as follows: the nodes can take two different roles, *aggregator* and regular *coalition node*. Among the nodes in the coalition, the aggregator acts as traffic relay, i.e., it forwards the traffic received from the other nodes via the short-range interface, using the cellular interface. Only one node can act as aggregator at a time. Fig. 3 exemplifies the general setup composed of 3 Watts2Share nodes using 3G and WiFi. Node 1 has the aggregator role and node 2 and 3 are regular coalition nodes. The aggregator role is shared among the nodes over time by performing role rotation.
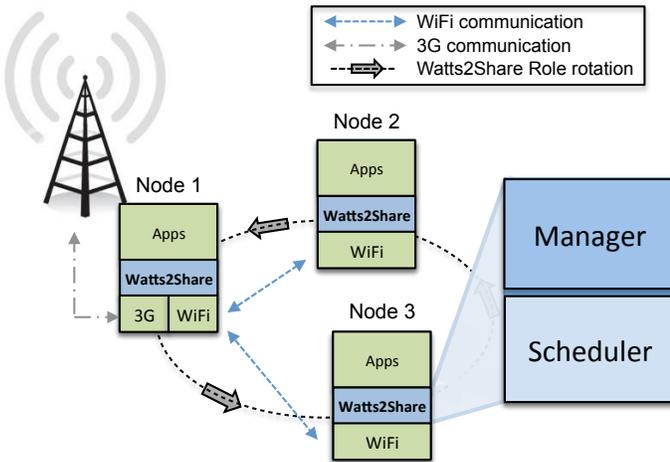
Fig. 3: Watts2Share architecture.

Fig. 3 shows that Watts2Share's design resides between the applications and the wireless interface, acting as a middleware component. Watts2Share is composed of 2 main components: scheduler and manager.

The scheduler is in charge of scheduling the traffic traces of the different nodes. In a regular coalition node, it simply schedules the application data using the short-range interface. In the aggregator, the scheduler performs two different tasks: (1) it takes care of the coalition nodes' traffic via the short-range interface, and (2) sends and receives the coalition nodes' and its own cellular data traffic. Since the aggregator acts as the main 3G sender and receiver and coordinates the short-range communication, the other coalition nodes are able to use power saving mechanisms when sending or receiving data reducing their energy consumption.

The manager takes care of the coalition creation and the aforementioned role rotation. The main purpose of the role rotation is to balance the energy consumption of the aggregator among all the coalition nodes. The manager thus performs the aggregator election (in a distributed manner) before switching roles.

When evaluating this architecture in section IV, we employ an instance of this architecture based on 3G-WiFi as a proof of concept. However, the approach is more general and can be applied to other (high/low energy) radio contexts. Our prototype simulates a simple scheduler by aggregating the traffic received from the short-range interface. This merges the traffic traces from all the nodes and bridges the short-range and cellular interface to perform traffic forwarding. Regarding role rotation, it implements a simple role rotation scheme where the aggregator role rotates every T_R in a round-robin fashion to balance the energy consumption among the current coalition nodes. The modular approach in the architecture allows plugging in more sophisticated decision algorithms based on energy levels or utility functions.

A. Energy consumption of Watts2Share

In this section we define metrics for the energy consumption of the two scenarios: (1) all nodes connected via their cellular interface in a standard scenario (E_S), and (2) in the Watts2Share scenario (E_W).

Consider a set of nodes $N = \{n_1, ..., n_k\}$ for both scenarios. The cumulative energy consumption of the standard scenario is the sum of the cellular energy consumption of each node,

$$E_S = \sum_{i=1}^{k} E_i^c \tag{2}$$

where E_i^c is the energy consumed by n_i using the cellular interface.

For Watts2Share, the nodes can adopt two roles, coalition node and aggregator. We define E_A to be the energy consumed by the aggregator and E_i^w the energy consumed by coalition node i using the WiFi interface. E_N is the sum of the energy of all the coalition nodes. Therefore, the energy consumption of the Watts2Share system is

$$E_W = E_A + E_N = E_A + \sum_{[i \in \{1,k\}, i \neq A]} E_i^w \tag{3}$$

The energy consumption of the aggregator is the sum of the cellular and WiFi consumption $E_A = E_A^c + E_A^w$, which depends on its own traffic and the traffic generated by the other nodes. The aggregator consumes energy for sending (own and coalition) uplink data in an aggregated manner, as well receiving corresponding downlink packets. Moreover, it has the extra cost of using the short-range interface for receiving/relaying data from/to the coalition nodes. We define the WiFi cost of acting as the aggregator as having the short-range interface always on in a high state, implying a worst case assumption. However, recent works [15] have shown that a software access point can sleep up to 60% of the time, which would reduce this cost for the aggregator.

IV. Evaluation

This section describes the experimental settings and the quantitative evaluation of the basic Watts2Share architecture.

A. Experimental settings and methodology

In order to define evaluation scenarios that are within the scope of the paper we need (1) to have an idea what the maximum bandwidth in the cellular network that we are using is, and (2) randomly select real data traffic in different nodes so that the assumption under section III (Equation 1) is valid. That is, the real data traffic traces should collectively use less than a proportion (ϕ) of the maximum available bandwidth. To determine the maximum 3G bandwidth at this particular location for an operator, we performed several downloads/uploads at different times of the day during a week. The maximum registered data rate was 6.23 Mbps and 1.32 Mbps for downlink and uplink respectively. The downlink bounds the maximum bandwidth of the system since most cellular data is downlink [16].

For our experiments, we set the threshold for the maximum available bandwidth to be a fraction $\phi = 0.85$ of the maximum registered. We measure and compute the maximum bandwidth during intervals of 50 ms, which provides higher granularity than measuring the maximum bandwidth over 1 second. The resulting limit is 264 kbits per 50 ms slot (or 5.3 Mbps).

We collect real packet level data traces of three different users using Android smartphones connected to a 3G network over the period of two weeks. The traces were collected running tcpdump in the Android devices. Each user performed her regular usage of the smartphone in terms of used applications (the users installed the applications that they use) and interaction with the smartphone. The traces contain both traffic created by the system applications and the user installed applications. The data traces contain traffic from commonly used applications like WhatsApp, Skype, Facebook, Angry Birds or Spotify that capture different real traffic patterns (the different users employ different applications). For example, a 5 minute fragment of a trace can contain the user interacting with Facebook for 1 minute, sending two instant messages using WhatsApp 2 minutes after, plus the rest of applications and system services running in background that also transfer data [7]. Then, we randomly select fragments of the traces that fall in the scope of this Watts2Share evaluation, i.e., collectively use up to 85% of the maximum bandwidth at any time.

The transmission energy consumption is calculated using EnergyBox. Given the 3G network parameters specified at operator level, EnergyBox derives the 3G states of the UE employing trace-based iterative packet-driven simulation. The total energy consumption is calculated by associating the UE specific power levels with the emulated intervals in each state, and integrating them over time. For WiFi, it simulates the the adaptive power save mode mechanism defined at the driver level of a WiFi device. EnergyBox has been evaluated against physical energy consumption measurements showing an average accuracy of 98% [17]. For the settings of the EnergyBox we used the 3G parameters of a real operator and the adaptive PSM parameters for the Galaxy SII driver. We used the power values previously observed in our measurement study in section II.

More specifically, for 3G we set T_1=5.1 s and T_2=12 s and the RLC buffer thresholds for FACH-DCH = 515 bytes. The state machine directly moves to DCH for any data transmission. We disabled the low activity mechanism described in section II and we did not use Fast dormancy. The power levels used were 612, 416, and 17 mW for DCH, FACH, and PCH respectively.

For WiFi we set $Up = 1$, $Down = 1$ and $\delta = 220$ ms. The power level for the states were: 12, 240, and 400 mW for PSM, CAM, and HighCAM respectively. The aggregator power level due to interacting with the WiFi coalition nodes was set to 346 mW based on our physical power measurements of the WiFi interface of a software access point (SoftAP). We run the SoftAP functionality in different devices (Sony Xperia Arc and Galaxy Nexus) and perform the measurements as described in section II. The measured average power values were 346 and 238 mW for the Xperia Arc and Galaxy Nexus respectively. Then, we selected the highest average power value corresponding to the worst case (Xperia Arc) for our tests.

For the following results, we calculate the energy consumption of the standard scenario (E_S) and the Watts2Share scenario (E_W) as specified in section III-A.

B. Watts2Share with two nodes

We start the evaluation by illustrating the most basic Watts2Share scenario composed of 2 nodes against the scenario where both nodes use their 3G interface (standard scenario). One node uses Google Translate whereas the other one checks the weather. The purpose of this study is to get a qualitative sense for the three different components of Watts2Share compared to the standard scenario.

The upper part of Fig. 4 shows the aggregate energy consumption of both nodes using their 3G and the individual consumptions. Node 2 stays in DCH state for most of the time, moving to FACH two times: 17 s and 31 s. Node 1 has similar behaviour, spending more time in DCH than node 2.

The lower part of Fig. 4 shows the aggregate energy consumption of two Watt2Share nodes. Node 1 is the aggregator and its energy consumption is composed of the short-range interface (aggregator role) plus forwarding both nodes' traffic via 3G. Node 2 is a coalition node that uses adaptive PSM in WiFi. The aggregator role results in node 2 moving to CAM only during packet transmissions. While not shown here, the choice of the aggregator in this scenario is not significant to our point (a similar chart would be shown in a reverse setting for the Watts2Share part).

The energy consumption of the standard scenario is 42.05 Joules, whereas Watts2Share consumes 32.29 J. Despite the high energy consumption of the aggregator, Watts2Share saves 23% of the energy for the same traffic. The aggregator consumes 44% more than what it would consume in the standard scenario. The standard scenario is expensive in terms
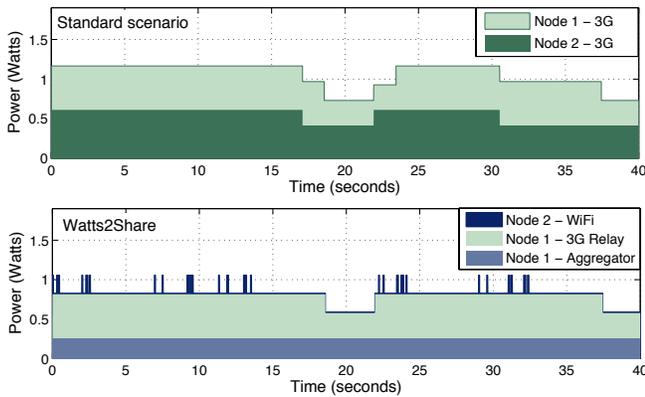
Fig. 4: Energy consumption of two 3G nodes compared to the same nodes using Watts2Share.

of energy due to the high energy consumption of 3G even for small amounts of traffic, where the energy tails become an important overhead. Instead, Watts2Share employs traffic consolidation by aggregating the traffic of both nodes into a single 3G stream, getting closer to the maximum traffic load where 3G becomes more energy-efficient. It is worth noting the coalition node spends 93% of the time in PSM, whereas in the standard scenario it does not even downswitch to the PCH state. These reasons balance the high cost of the aggregator leading to overall energy savings.

C. Increasing number of nodes

In this section we show that Watts2Share's traffic consolidation technique benefits from increasing the number of coalition nodes as long as the assumptions on the scope (below max 3G bandwidth) holds. We study the energy savings of Watts2Share when increasing the number of nodes.

We use the previously described real packet traces in fragments of 5 minutes (chunks) to emulate traffic originated at several nodes (up to 8) and obtain the energy savings. Then, we repeat every experiment 10 times using different data chunks, and present the mean and the variance. The results in Fig. 5 are normalised to the average energy consumption for the 2 node case of the standard scenario.
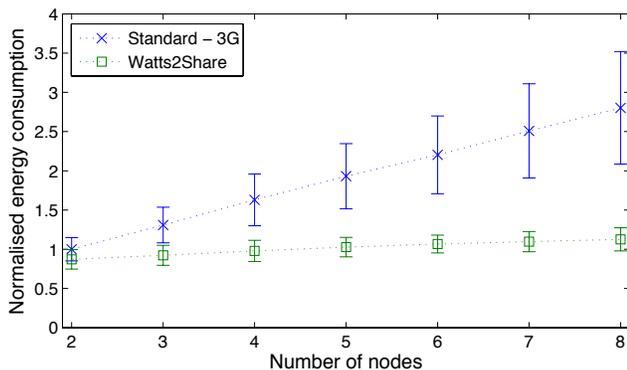
Fig. 5: Energy savings in relation to number of nodes.

Fig. 5 shows that in the standard scenario, the energy consumption steadily increases with more nodes. Each node contributes with a high energy consumption due to switching to DCH and FACH even though the data traffic is low. The high variance describes the high energy consumption difference for different traffic chunks that may include different applications with different data requirements.

Fig. 5 shows that the energy consumption of Watts2Share barely increases when adding more coalition nodes and the average energy savings are significantly increased from 12% for the 2 nodes Watts2Share, to 59% for the 8 nodes Watts2Share. Note that the main contributor to the energy consumption of Watts2Share, the aggregator, is already part of the two nodes Watts2Share, thereby achieving an almost linear saving after this number of nodes (and up to the maximum 3G load). The variance is lower than in the standard scenario since the WiFi consumption difference for different traces is smaller than for 3G.

The aggregator spends almost 100% of the time in DCH even with small number of nodes where the bandwidth load is far under the maximum bandwidth. Therefore, the more nodes we add, the closer we get to fully utilising the available bandwidth for the same energy consumption, thus achieving higher energy proportionality. However, the number of nodes is bounded by their aggregated bandwidth claim and the current available bandwidth at the aggregator node. Moreover, the greater the energy consumption difference between 3G and WiFi for a coalition node, the higher the aggregate energy savings become.

D. 3G signalling

In this section we illustrate that as well as saving energy for the handsets, Watts2Share has benefits for the cellular operator. Apart from reducing the 3G interferences due to reducing the number of connected nodes, we show that Watts2Share reduces the amount of signalling messages introduced by the RRC state transitions. In the normal scenario, every node performs the corresponding signalling with the RNC for allocating the physical channels for every state transition. For example, a single FACH-DCH state transition is associated with 10 message signalling transmissions [18]. When the traffic load is not high, state transitions are frequent, creating great signalling overhead.

Our approach to show the reduced signalling is as follows: for the experiments performed in section IV-C we count the number of state transitions recorded by the 3G EnergyBox. We report the average number of state transitions over the 10 experiments as well as the variance. The results of Fig. 6 are normalised by dividing the number of state transitions by the case of 2 nodes in the standard 3G scenario.

Fig. 6 shows that, as expected, the number of state transitions of the nodes increases with more nodes connected directly to the 3G network. The number of transitions depends on the data pattern and can vary a lot from node to node (as vivid in the variance). However, in Watts2Share, the number of state transitions is fairly constant since the aggregator
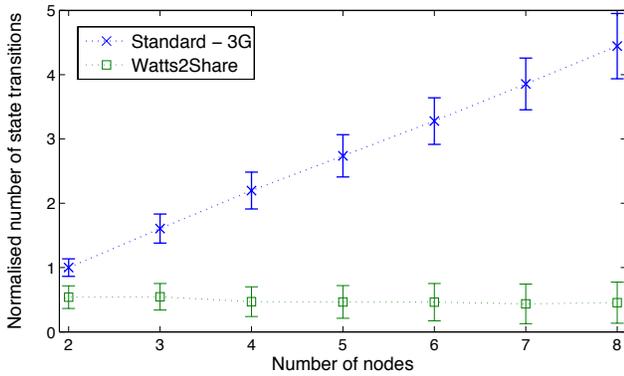
Fig. 6: State transition reduction of Watt2Share.

spends most of the time in DCH state. When the number of nodes is more than 3, the only state transitions experienced by the aggregator are between DCH and FACH and rarely downswitching to PCH. This leads to an average reduction of state transitions ranging from 46% to 89% depending on the number of nodes.

E. System lifetime

The energy burden of the aggregator should be balanced among the coalition nodes with a role rotation scheme. In order to study the impact of the rotation scheme, and its variation depending on frequency of rotation, we performed a number of 4 node experiments with T_R set to 5 minutes for the round robin rotations. We plot the energy drain for each node as well as the system lifetime. Let's define system lifetime as the interval during which all four nodes are operational. The nodes start with different energy levels randomly set to simulate a real-life scenario and drain their energy during the experiment time depending on their transmissions (x and y axes in Fig. 7 respectively). The starting energy of the nodes is normalised to the energy of node 1.
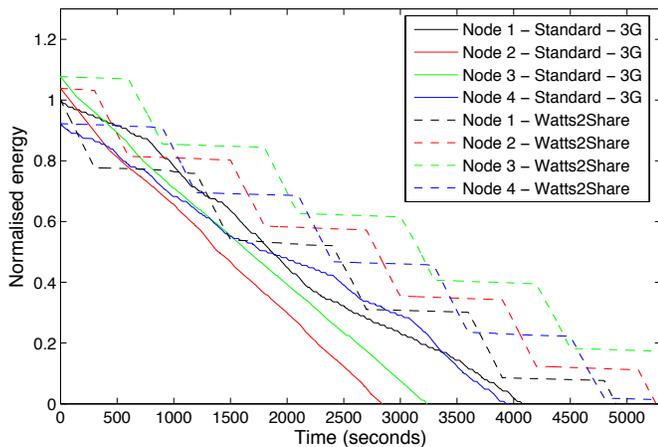
Fig. 7: System lifetime of Watts2Share.

Fig. 7 shows that the system lifetime within the standard scenario is bounded by node 2, which runs out of energy after

2816 seconds. Node 4 is the last node running out of energy after 4086 seconds. Watts2Share extends the system lifetime by 73%: node 1 is the first node running out of battery in Watts2Share after 4890 seconds. The other nodes have even longer lifetime.

Fig. 7 also depicts the impact of being an aggregator. Node 1 starts being the aggregator for 5 minutes, followed by node 2. The aggregator consumes much more energy than the coalition nodes which is shown by a steeper drop in energy. The fact that the aggregator spends most of its time in the DCH state causes the energy to drop sharply. When node 1 switches from aggregator to coalition node, it consumes only using WiFi, which makes the energy drain much smaller. This behaviour is similar for all nodes.

Note that different nodes experience different energy savings. For example, node 1 is not the one with the least energy in the Watts2Share scenario, but it runs out of energy first. The value of T_R greatly impacts the energy balancing. Next section discusses how to set the value of T_R to balance the energy consumption among the coalition nodes.

F. Role rotation overhead

There is a trade-off when choosing the T_R parameter: shorter rotation times allow a fairer distribution of the co-ordination overhead, but also increase the overall experienced overhead since the rotation itself has an energy footprint (CPU as well as messaging).

In order to quantify the switching overhead in terms of message exchanges needed, we would need to analyse the algorithms for rotation. As a first step, we measure the energy consumption and time taken to switch on/off the SoftAP functionality in a modern smartphone (Sony Xperia Arc). Activating the SoftAP takes 2.98 seconds and consumes 0.23 Joules.

The energy overhead is defined as the collective energy spent by all nodes for each rotation of the aggregator role. The time overhead relates to the time that the nodes collectively spend for the purpose of rotating the aggregator role. When a rotation occurs, two nodes switch their role: the aggregator and a coalition node. These two nodes experience a time overhead in terms of disruption time, when they cannot communicate since they are rotating their role. The other coalition nodes perform a handover to the new aggregator.

Fig. 8 shows the impact of selecting T_R on the total activation energy and time overhead for the scenario of 4-node Watts2Share from section IV-E. We select T_R values that range from 1 minute to 15 minutes. The time and energy overhead values are normalised to the duration of the experiment (5300s) and the node with the lowest initial energy (1150 Joules), respectively.

Fig. 8 shows that for short T_R the time and energy overheads are 5 and 3.5% respectively. However, when increasing the T_R to higher values the overhead experiences a visible decrease, settling around 1% when T_R is 5 minutes, the value used in earlier sections. The time overhead seems more significant than the energy overhead since the handsets seem to be
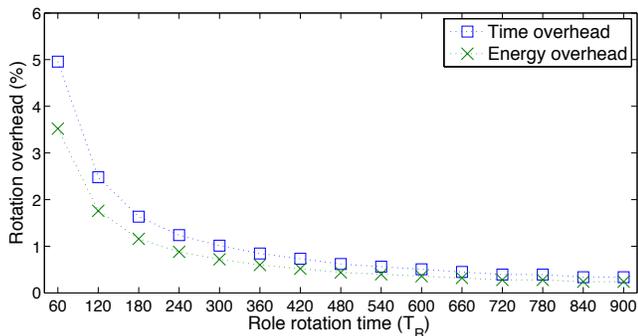
Fig. 8: Role rotation overhead.

more optimised towards low energy SoftAP operation than performance degradation from frequent switching.

Selecting a T_R value depends also on the period of time that the nodes collaborate using Watts2Share. When the nodes form Watts2Share coalitions for longer intervals, greater T_R values can be adopted to further decrease the overhead. However, for short periods, a shorter T_R is suggested for balancing the energy consumption over the nodes.

G. Other coalition overheads

So far we focused on showing the benefits of Watts2Share in terms of energy savings. There are of course some costs associated with the scheme including coalition formation energy costs and role rotation energy cost. While the role rotation cost for an aggregator swap was presented in section IV-F, the other element needs more extensive studies in a physical testbed.

A preliminary implementation of Watts2Share in a modern Android device (LG Nexus 4) shows the radio energy cost for a single coalition formation at the client and the aggregator respectively [28]. The coalition formation is divided in three stages: SoftAP activation for the aggregator, parameter exchange (e.g., network configuration), and connection. The measurements indicate 0.6 Joules over 1.6 seconds for SoftAP activation for the aggregator, while the client is in idle consuming 0.02 Joules. This is followed by a parameter exchange between the client and the aggregator which costs 0.3 Joules over 0.8 seconds at the client, and 0.26 Joules for the aggregator[1]. Finally, 0.5 Joules over 2 seconds are consumed for the connection at the client, and 0.7 Joules for each client connection at the aggregator. The total process is 4.2 seconds and costs 0.82 and 1.56 Joules to the client and the aggregator respectively.

These measurements indicate a promising potential with a significant net saving by Watts2Share in underutilised scenarios.

V. RELATED WORK

Forming ad hoc configurations as opposed to using an existing infrastructure is a massively studied topic [19], [20]. Energy consumption is known to be the Achilles' heel of

---

[1]Since the aggregator has the SoftAP activated, this consumes almost as much as the client.

ad hoc communication [21], [22]. The aim of Watts2Share is to reduce the overall footprint of a group by avoiding the continuous listening mode of the WiFi connected nodes, and using an aggregator to reach the cellular infrastructure.

From here on, we categorise the main body of related work into two areas: energy consumption studies and energy saving techniques for wireless networking. Some studies and proposals consider infrastructure-centric approaches for 3G [23] and WiFi [13], whereas our focus is on studying the energy consumption at the user side.

Energy saving techniques: we further divide the related works into energy saving techniques at a single handset level or considering a combination of handsets. Within one handset, Balasubramanian et al. [9] introduce the concept of Delay Tolerant Applications (DTA) via the intuition that to reduce energy consumption some applications can defer their transmissions. TailEnder performs bursty transmissions attempting to reduce the number of state transitions. Our previous work [7] refines the above work by reducing unnecessary state transitions and energy consumption. Several works employ Fast Dormancy to reduce the energy consumption of cellular traffic [24], [25]. Schulman et al. [26] explore the impact of the network coverage on the energy consumption.

For WiFi, some works [14], [27], [29] focus on optimising the PSM operation for different traffic loads. DozyAP [15] focuses on reducing the energy consumption of WiFi tethering and shows that the WiFi interface of a SoftAP can sleep up to 88% of the time. Camps-Mur et al. [30] study and propose power management protocols to reduce the energy consumption of SoftAPs using WiFi Direct.

The above works are orthogonal to Watts2Share and they can be implemented either in the aggregator to reduce 3G energy consumption, or for the coalition nodes in the WiFi case to achieve higher energy savings.

When considering energy saving techniques for multiple nodes, Cool-Tether [31] is a WiFi hotspot focused on serving Web pages using the cellular connection from smartphones. A cloud-based server fetches and reduces the data of the Web page requests and sends it in a single burst to a certain number of smartphones that act as forwarders speeding up the transfer. Their work offloads the energy burden of the WiFi AP from mobile phones to a laptop using reverse-infrastructure mode. Our work maximises the system lifetime by only considering the limited energy reserve of mobile nodes with using no cloud-server assistance. Moreover, our work is evaluated using real traffic generated by a set of applications and real user interactions.

Yoo et al. [32] present a cooperative approach to reduce the energy consumption of WiFi creating Bluetooth Personal Area Networks (PAN). Watts2Share has a similar aspiration but exploits the knowledge of how the aggregated data transmission taxes the recurring (rotating) relay node's energy in a 3G context.

Perruci et. al [33] propose a cooperative mobile web browsing approach combining Bluetooth and cellular networks to decrease Web download time. In their architecture, two nodes

(master-slave) are interested in the same data, and therefore the nodes download parts of the Web and share it afterwards. In the same context, Perrucci et al. [34] also introduce a TDMA based MAC layer scheme to improve the energy efficiency of cooperation in wireless communication. Watts2Share differs in that each node is interested in its own transmissions. We build on standard protocols for transmission and adds a thin layer below the applications.

Asadi et al. [5] present a cooperative uplink packet forwarding mechanism employing a bi-radio (WiFi-LTE) scenario assuming fully utilised links. In comparison, our work focuses on a underutilised scenario and considers uplink and downlink real user usage traces for the evaluation.

Lei et al. [35] complement our work by analysing the business models in the case of operator controlled peer-to-peer communication in combination with Long Term Evolution. Their work provides cooperation using different incentives. We believe that our suggested operator and energy-aware relays are a good example of such incentives. Guo et al. [36] propose a cooperative relay service within a WiFi infrastructure. The terminals that anyway need to be connected to the AP due to their high data rate to forward, carry also data for low data rate nodes (and get an incentive). Our work focuses on energy gains due to an energy-aware use of the ongoing 3G connection.

Compared to the wireless sensor networking literature [37]–[39], in Watts2Share there is no common goal for the whole network (e.g., data collection), and the traffic load for the individual nodes is different and unpredictable.

To sum up, in comparison to previous works, we evaluate Watts2Share using real user traces of common applications capturing the energy impact of the traffic pattern in a hybrid 3G-WiFi scenario.

Energy consumption studies: A detailed study of the tail energy overhead in 2G/3G cellular networks is performed by Qian et al. [40] on user data traces retrieved from a network operator in 2009. The work is extended [1] pointing out how different applications inefficiently utilise the radio resources due to their data pattern. A recent study by Pathak et al. [2] focuses on profiling energy consumption of smartphone applications using a system-call-based power modelling for smartphones. Rice et al. [8] measure the WiFi energy consumption in a variety of smartphones and Friedman et al. [41] present a comparative study between WiFi and Bluetooth.

VI. CONCLUSIONS AND FUTURE WORK

In the quest for energy efficiency in wireless communication, energy proportionality plays an important role. A typical 3G user does not use its available 3G bandwidth up to the maximum all the time. Contrary to a naive expectation, even when the maximum available bandwidth is underutilised, the user still experiences high energy consumption and short battery lifetimes (thereby low quality of experience).

In our work we propose Watts2Share, an architecture for energy-efficient wireless communication that reduces the energy consumption compared with a pure 3G scenario by performing traffic consolidation on a single link. The nodes form coalitions employing WiFi where a member retransmits other nodes' traffic. We show that Watts2Share can benefit all the end users and reduces overall energy by reducing the amount of signalling. Watts2Share further extends the lifetime of all coalition nodes by using a role rotation scheme.

The architecture of Watts2Share is not limited to 3G and WiFi. We believe that our energy-aware traffic consolidation approach is applicable to other technologies as well, such as LTE and WiFi Direct or 3G and Bluetooth. The higher data rates and energy consumption of LTE would allow a higher number of coalition nodes and potentially greater energy savings.

Our current work includes implementation of the scheme using native code in an Android phone, and optimising the role rotation mechanism by incorporating the EnergyBox into the manager. The use of EnergyBox will enable energy quantification towards energy-aware role rotation decisions. Future work includes building on the gains that Watts2Share achieves in this paper through seamless adaptation of coalitions during periods of overload and underload. Watts2Share needs to be extended to a nonuniform scenario where some flows cannot be disrupted without a quality of service penalty or not all nodes are equally eligible as an aggregator. More work is needed to extend the scheme and the physical testbed studies for these scenarios.

REFERENCES

[1] F. Qian, Z. Wang, A. Gerber, Z. Mao, S. Sen, and O. Spatscheck, "Profiling Resource Usage for Mobile Applications: a Cross-Layer Approach," in *Proceedings of the 9th international conference on Mobile systems, applications, and services*, MobiSys '11. ACM, 2011, pp. 321–334.

[2] A. Pathak, Y. C. Hu, and M. Zhang, "Where is the energy spent inside my app?: Fine Grained Energy Accounting on Smartphones with Eprof," in *Proceedings of the 7th ACM european conference on Computer Systems*, EuroSys '12. ACM, 2012, pp. 29–42.

[3] B. Han, P. Hui, V. Kumar, M. Marathe, J. Shao, and A. Srinivasan, "Mobile Data Offloading Through Opportunistic Communications and Social Participation," *IEEE Transactions on Mobile Computing,*, vol. 11, no. 5, pp. 821 –834, May 2012.

[4] K. J. O'Brien. (2012, Sep.) From London, a lesson in the benefits of free Wi-Fi. [Online]. Available: http://www.nytimes.com/2012/09/03/technology/from-london-a-lesson-in-the-benefits-of-free-wi-fi.html

[5] A. Asadi and V. Mancuso, "Energy Efficient Opportunistic Uplink Packet Forwarding in Hybrid Wireless Networks," in *Proceedings of the 4th International Conference on Future Energy Systems*, e-Energy '13. ACM, 2013, pp. 261–262.

[6] U. Paul, A. Subramanian, M. Buddhikot, and S. Das, "Understanding Traffic Dynamics in Cellular Data Networks," in *IEEE INFOCOM*, 2011.

[7] E. J. Vergara and S. Nadjm-Tehrani, "Energy-Aware Cross-Layer Burst Buffering for Wireless Communication," in *Proceedings of the 3rd International Conference on Future Energy Systems: Where Energy, Computing and Communication Meet*, e-Energy '12. ACM, 2012.

[8] A. Rice and S. Hay, "Measuring Mobile Phone Energy Consumption for 802.11 Wireless Networking," *Journal Pervasive Mobile Computing*, vol. 6, 2010.

[9] N. Balasubramanian, A. Balasubramanian, and A. Venkataramani, "Energy Consumption in Mobile Phones: A Measurement Study and Implications for Network Applications," in *Proceedings of the International Measurement Conference*, IMC, ACM 2009.

[10] H. Holma and A. Toskala, *WCDMA for UMTS: HSPA Evolution and LTE*, ser. Wiley Online Library: Books. John Wiley & Sons, 2010.

[11] R. Krashinsky and H. Balakrishnan, "Minimizing Energy for Wireless Web Access with Bounded Slowdown," in *Proceedings of the 8th Annual International Conference on Mobile Computing and Networking*, MobiCom '02. ACM, 2002.

[12] A. J. Pyles, Z. Ren, G. Zhou, and X. Liu, "SiFi: Exploiting VoIP Silence for WiFi Energy Savings in Smart Phones," in *Proceedings of the 13th International Conference on Ubiquitous Computing*, UbiComp '11. ACM, 2011.

[13] J. Manweiler and R. Roy Choudhury, "Avoiding the Rush Hours: Wifi Energy Management via Traffic Isolation," *IEEE Transactions on Mobile Computing,*, vol. 11, no. 5, pp. 739–752, may 2012.

[14] A. J. Pyles, X. Qi, G. Zhou, M. Keally, and X. Liu, "SAPSM: Smart Adaptive 802.11 PSM for Smartphones," in *Proceedings of the 14th International Conference on Ubiquitous Computing*, UbiComp '12. ACM, 2012.

[15] H. Han, Y. Liu, G. Shen, Y. Zhang, and Q. Li, "DozyAP: Power-Efficient Wi-Fi Tethering," in *Proceedings of the 10th International Conference on Mobile Systems, Applications, and Services*, MobiSys '12. ACM, 2012, pp. 421–434.

[16] H. Falaki, D. Lymberopoulos, R. Mahajan, S. Kandula, and D. Estrin, "A First Look at Traffic on Smartphones," in *Proceedings of the 10th Annual Conference on Internet Measurement*, IMC '10. ACM, 2010, pp. 281–287.

[17] E. J. Vergara and S. Nadjm-Tehrani, "EnergyBox: A Trace-Driven Tool for Data Transmission Energy Consumption Studies," in *Proceedings of the International Conference on Energy Efficiency in Large Scale Distributed Systems (EE-LSDS 2013)*, Lecture Notes in Computer Science. Springer, 2013.

[18] L. Qian, E. W. W. Chan, P. P. C. Lee, and C. He, "Characterization of 3G Control-Plane Signaling Overhead from a Data-Plane Perspective," in *Proceedings of the 15th ACM international conference on Modeling, analysis and simulation of wireless and mobile systems*, MSWiM '12. ACM, 2012.

[19] I. Chlamtac, M. Conti, and J. J.-N. Liu, "Mobile Ad Hoc Networking: Imperatives and Challenges," *Ad Hoc Networks*, vol. 1, no. 1, pp. 13 – 64, 2003.

[20] H. Luo, R. Ramjee, P. Sinha, L. E. Li, and S. Lu, "UCAN: a Unified Cellular and Ad-Hoc Network Architecture," in *Proceedings of the 9th Annual International Conference on Mobile Computing and Networking*, MobiCom '03. ACM, 2003.

[21] E. J. Vergara, S. Nadjm-Tehrani, M. Asplund, and U. Zurutuza, "Resource Footprint of a Manycast Protocol Implementation on Multiple Mobile Platforms," in *5th International Conference on Next Generation Mobile Applications, Services and Technologies (NGMAST)*, Sept. 2011.

[22] S. Trifunovic, B. Distl, D. Schatzmann, and F. Legendre, "WiFi-Opp: Ad-Hoc-less Opportunistic Networking," in *Proceedings of the 6th ACM Workshop on Challenged Networks*, CHANTS '11. ACM, 2011.

[23] I. Humar, X. Ge, L. Xiang, M. Jo, M. Chen, and J. Zhang, "Rethinking Energy Efficiency Models of Cellular Networks with Embodied Energy," *Network, IEEE*, vol. 25, no. 2, pp. 40 –49, 2011.

[24] F. Qian, Z. Wang, A. Gerber, Z. Mao, S. Sen, and O. Spatscheck, "TOP: Tail Optimization Protocol for Cellular Radio Resource Allocation," in *Proceedings of the 18th IEEE International Conference on Network Protocols (ICNP),* , 2010.

[25] P. Athivarapu, R. Bhagwan, S. Guha, V. Navda, R. Ramjee, D. Arora, V. Padmanabhan, and G. Varghese, "RadioJockey: Mining Program Execution to Optimize Cellular Radio Usage," in *Proceedings of the 18th Annual International Conference on Mobile Computing and Networking*, Mobicom'12 2012.

[26] A. Schulman, V. Navda, R. Ramjee, N. Spring, P. Deshpande, C. Grunewald, K. Jain, and V. N. Padmanabhan, "Bartendr: a Practical Approach to Energy-Aware Cellular Data Scheduling," in *Proceedings of the 16th Annual International Conference on Mobile Computing and Networking,* Mobicom'10, 2010.

[27] E. Tan, L. Guo, S. Chen, and X. Zhang, "PSM-Throttling: Minimizing Energy Consumption for Bulk Data Communications in WLANS," in *Proceedings of the IEEE International Conference on Network Protocols,* (ICNP), IEEE 2007.

[28] U. Moreno, "Energy Consumption Studies for 3G Traffic Consolidation on Android using WiFi and Bluetooth," Master Thesis, Linköping University, 2013.

[29] Y. He and R. Yuan, "A Novel Scheduled Power Saving Mechanism for 802.11 Wireless LANs," *IEEE Transactions on Mobile Computing,*, vol. 8, no. 10, pp. 1368–1383, 2009.

[30] D. Camps-Mur, X. Pérez-Costa, and S. Sallent-Ribes, "Designing Energy Efficient Access Points with Wi-Fi Direct," *Comput. Netw.*, vol. 55, no. 13, Elsevier, 2011.

[31] A. Sharma, V. Navda, R. Ramjee, V. N. Padmanabhan, and E. M. Belding, "Cool-Tether: Energy Efficient On-the-fly WiFi Hot-spots Using Mobile Phones," in *Proceedings of the 5th International Conference on Emerging Networking Experiments and Technologies*, CoNEXT '09. ACM, 2009.

[32] J. W. Yoo and K. H. Park, "A Cooperative Clustering Protocol for Energy Saving of Mobile Devices with WLAN and Bluetooth Interfaces," *IEEE Transactions on Mobile Computing*, vol. 10, 2011.

[33] G. P. Perrucci, F. H. P. Fitzek, Q. Zhang, and M. D. Katz, "Cooperative Mobile Web Browsing," *EURASIP Journal on Wireless Communications and Networking*, vol. 2009, April 2009.

[34] G. P. Perrucci, P. Anggraeni, S. Wardana, F. H. P. Fitzek, and M. Katz, "Bio-Inspired Energy-Aware Protocol Design for Cooperative Wireless Networks," *International Journal of Autonomous and Adaptive Communications Systems (IJAACS)*, 2009.

[35] L. Lei, Z. Zhong, C. Lin, and X. Shen, "Operator Controlled Device-to-Device Communications in lte-advanced networks," *IEEE Wireless Communications,* , vol. 19, no. 3, pp. 96–104, June 2012.

[36] L. Guo, X. Ding, H. Wang, Q. Li, S. Chen, and X. Zhang, "Cooperative Relay Service in a Wireless LAN," *IEEE Journal on Selected Areas in Communications*, vol. 25, no. 2, pp. 355–368, 2007.

[37] W. R. Heinzelman, A. Chandrakasan, and H. Balakrishnan, "Energy-Efficient Communication Protocol for Wireless Microsensor Networks," in *Proceedings of the 33rd Hawaii International Conference on System Sciences - Volume 8*, HICSS '00. IEEE, 2000, pp. 8020–.

[38] G. Anastasi, M. Conti, M. D. Francesco, and A. Passarella, "Energy Conservation in Wireless Sensor Networks: A Survey," *Ad Hoc Networks*, vol. 7, no. 3, pp. 537 – 568, 2009.

[39] O. Boyinbode, H. Le, A. Mbogho, M. Takizawa, and R. Poliah, "A Survey on Clustering Algorithms for Wireless Sensor Networks," in *Proceedings of the 13th International Conference on Network-Based Information Systems (NBiS)*, Sept. 2010, pp. 358–364.

[40] F. Qian, Z. Wang, A. Gerber, Z. M. Mao, S. Sen, and O. Spatscheck, "Characterizing Radio Resource Allocation for 3G Networks," in *Proceedings of the 10th Annual Conference on Internet Measurement*, ser. IMC '10. ACM, 2010.

[41] R. Friedman, A. Kogan, and Y. Krivolapov, "On Power and Throughput Tradeoffs of WiFi and Bluetooth in Smartphones," in *INFOCOM, Proceedings IEEE*, April 2011.