

Enabling always on service discovery: Wi-Fi Neighbor Awareness Networking (NAN)

Journal:	<i>IEEE Wireless Communications Magazine</i>
Manuscript ID:	WCM-14-00271
Topic or Series:	Open Call Article
Date Submitted by the Author:	01-Oct-2014
Complete List of Authors:	Camps-Mur, Daniel; NEC Network Laboratories, Garcia-Villegas, Eduard; UPC, Lopez-Aguilera, Elena; UPC, Loureiro, Paulo Lambert, Paul Raissinia, Ali
Key Words:	Wi-Fi NAN, Service Discovery, Low Power, Synchronization

SCHOLARONE™
Manuscripts

Pre-Review Only

Enabling always on service discovery: Wi-Fi Neighbor Awareness Networking (NAN)

Daniel Camps-Mur, Eduard Garcia-Villegas, Elena Lopez-Aguilera, Paulo Loureiro, Paul Lambert and Ali Raissinia

Abstract—There is an untapped potential in the Wi-Fi radios embedded in our smartphone and tablet devices. In this paper we introduce the Wi-Fi Neighbor Awareness Networking (NAN) technology, being standardized in the Wi-Fi Alliance, which leverages this potential by allowing handheld devices to continuously discover other interesting services and devices, while operating in the background in an energy efficient way. In addition, we present a thorough performance evaluation based on packet level simulations that illustrates the performance of Wi-Fi NAN to be expected in realistic scenarios.

Index Terms—Wi-Fi NAN, Service Discovery, Low power, Synchronization

I. INTRODUCTION

The Wi-Fi technology is currently embedded in most of the smartphone and tablet devices that people carry around while on the move. These embedded radios though are usually only operated when the user directly interacts with the device, for instance to access the Internet. Hence, there is a large spectrum of novel applications that could be devised if the Wi-Fi radios in our devices would continuously operate in the background discovering interesting devices or services on behalf of their users.

The biggest hurdle to a continuous background operation is the fact that current radio technologies, like Wi-Fi, are power hungry and do not allow a handheld device to keep its radio continuously active without heavily impacting battery life.

In order to address this challenge, in this paper we present the Wi-Fi Neighbor Awareness Networking (NAN) technology, currently being standardized in the Wi-Fi Alliance [1]. The underlying idea behind Wi-Fi NAN is to synchronize a set of potentially neighboring Wi-Fi devices, thus enabling an energy efficient low duty cycle operation, where devices wake up, listen, and broadcast small chunks of information within small Discovery Windows (DWs), as illustrated in the lower part of Figure 1. In Wi-Fi NAN, the achieved duty cycles guarantee a continuous operation of the technology in a smartphone or tablet device.

D. Camps-Mur (daniel.camps@i2cat.net) is with the i2CAT Foundation in Barcelona, Spain.

E. Garcia-Villegas (eduardg@entel.upc.es) and E. Lopez-Aguilera (elopez@entel.upc.es) are with Universitat Politcnica de Catalunya - BarcelonaTech (UPC) in Barcelona, Spain.

P. Loureiro (loureiro@neclab.eu) is with NEC Network Laboratories in Heidelberg, Germany.

P. Lambert (paul@marvell.com) is with Marvell Technology in Santa Clara, USA.

A. Raissinia (alirezar@qualcomm.com) is with Qualcomm in San Diego, USA.

The design of Wi-Fi NAN draws on previous work related to low duty cycle MAC protocols for Wireless Sensor Networks (WSNs), like [2] and [3]. However, the unique requirements spawning from current smartphone platforms and from the expected mobility patterns justify the need for the novel solutions devised in Wi-Fi NAN.

The paper is organized as follows. Section II introduces the architecture of Wi-Fi NAN. Section III contains a detailed description of the novel MAC aspects in NAN, and Section IV describes the way applications interact with the technology. Section V contains a detailed performance evaluation based on packet level simulations. Finally, section VI concludes the paper.

II. WI-FI NAN ARCHITECTURE

The Wi-Fi NAN is built upon the interaction of NAN devices grouped in clusters. Clusters are automatically created by nearby NAN devices that cooperate to synchronize to a common DW schedule. During that DW, all NAN devices participating in the cluster are allowed to exchange service frames describing or requesting a service. A NAN device is any Wi-Fi capable device supporting all required NAN protocol mechanisms. The NAN stack has two main components, the Discovery Engine (DE), providing basic Publish/Subscribe mechanisms to upper-layer services or applications, and the NAN Medium Access Control (MAC), responsible for the maintenance of NAN Clusters (creating, joining or merging clusters), for preserving synchronization in the NAN Cluster, and for providing transmit and receive services to the DE. Within a NAN cluster, a NAN device can operate under different roles which entail different responsibilities: Master or Non-Master. The upper part of Figure 1 illustrates the architecture of a NAN device.

III. THE NAN MAC

A. NAN Cluster Discovery

NAN devices discover NAN clusters through scanning in a particular pre-defined channel: channel 6 (2.437 GHz) in the 2.4 GHz band, channel 44 (5.220 GHz) in the 5 GHz lower band (5.150 – 5.250 GHz), channel 149 (5.745 GHz) in the 5 GHz upper band (5.725 – 5.825 GHz), and channel 149 if both 5 GHz upper and lower bands are allowed.

In order to allow NAN cluster discovery, each NAN device operating in Master role broadcasts a specific management frame called Discovery Beacon outside the DW (with an average transmission period of 100 ms), as illustrated in the lower part of Figure 1. In order to minimize

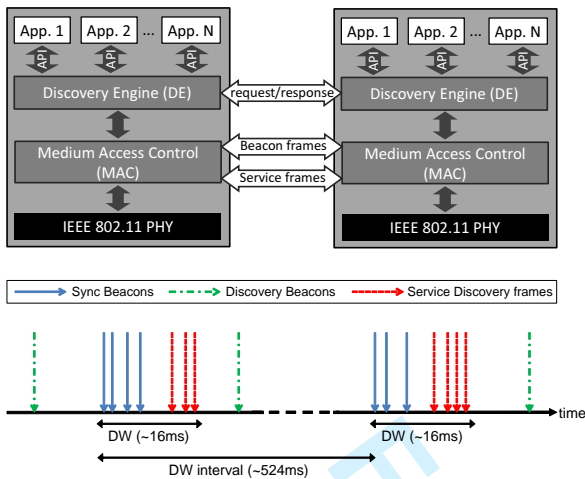


Fig. 1. NAN Architecture and operation in DW

the energy required to transmit Discovery Beacons, such a transmission is prioritized in front of other regular Wi-Fi transmissions, and it is skipped in case the intended transmission time overlaps with the DW of the corresponding NAN cluster. Thus, discovery of NAN clusters is achieved by having NAN devices passively scan for Discovery Beacons with a frequency decided by each implementation, for instance, according to the device's power budget.

The Discovery Beacon frame is based on the original IEEE 802.11 Beacon management frame format [4], modified to include NAN specific information such as the Cluster ID. The Cluster ID, included in all NAN protocol transmissions, is randomly chosen by the device starting the NAN cluster; thus, different IDs are generated for the identification of different clusters.

If, after the passive scan, a recently initiated NAN device does not detect any cluster, it can start a new NAN cluster. Instead, if more than one cluster were discovered, the device chooses the cluster to connect according to a selection method specified in the standard (cf. III-D).

B. NAN Device States

Two mechanisms are essential in a NAN cluster. First, NAN devices must be able to discover existing NAN clusters with a minimal power consumption. Second, NAN devices in a cluster must be able to synchronize their clocks in order to maintain their DWs aligned and be able to exchange service discovery frames. For these two purposes, NAN devices transmit Discovery (cf. III-A) and Synchronization Beacons (cf. III-C). A core responsibility of the NAN protocol consists in distributing the task of Beacon generation among the devices in a cluster. The mechanism designed allows devices to express a preference towards Beacon generation (e.g. devices without battery restrictions could express higher preferences) while, at the same time, fairness among devices with the same preference is achieved. For this purpose, the NAN protocol specifies a set of roles or states: Anchor Master, Master, Non-Master Sync or Non-Master Non-Sync.

To achieve synchronization, all devices in a cluster need to follow the same clock source. Thus, in NAN, the Anchor Master device is the main entity responsible for maintaining the synchronization used to align the DW for service discovery functions, and thus, all devices in the cluster follow the Anchor Master's time reference through their Time Synchronization Function (TSF). Even though a NAN cluster may temporarily have different Anchor Masters, the procedures of the NAN protocol ensure that a NAN cluster always converges to having only one Anchor Master.

NAN Devices operating in Master role are responsible for propagating both synchronization and discovery information of the cluster by sending Synchronization and Discovery Beacon frames respectively.

Devices in Non-Master Sync role participate in the propagation of Synchronization Beacon frames but are relieved from transmitting Discovery Beacon frames. Indeed, the need for the Non-Master Sync state stems from the fact that, due to their location within the cluster, some devices must be eventually forced to forward synchronization information in order to keep the cluster synchronized, despite expressing a lower preference. Finally, devices in Non-Master Non-Sync role are relieved from the task of propagating both Synchronization and Discovery Beacons. In addition, Non-Master Non-Sync devices need not be awake during all DWs and can therefore benefit from larger energy savings.

The assignment of states to devices by the NAN mechanisms is intended to maximize dissemination of Synchronization Beacons and the range of Discovery Beacons throughout the whole cluster, while employing a number of Master and Non-Master Sync devices as low as possible, hence minimizing the number of transmitted Beacon frames. In consequence, For these reasons, the mechanism used to arbitrate the state transitions of a device (which determines how devices share the burden of Beacon transmission) is one of the key components of NAN.

In order to fairly distribute energy savings, each NAN device manages a *NAN Master Rank* value, which is ensured to be unique while balancing preference and fairness; the NAN device with the highest Master Rank in the cluster becomes the Anchor Master. The Master Rank value is computed as a function of three components: a Master Preference value (that may change in time), a Random Factor value (that is periodically updated) and the device's MAC address. A higher value of the Master Preference means a device's higher preference to serve as a Master. In this sense, devices with less stringent battery requirements are recommended to choose larger Master Preference values. The Random Factor guarantees that devices with equal Master Preference will have equal chance of assuming the Master role.

Initially, when a NAN device joins a cluster, it assumes the Master role. The device transitions its role to Non-Master Sync when it becomes aware of the presence of one or more Master devices in proximity that manage a higher Master Rank. On the other hand, a Non-Master device (either Sync or non-Sync) assumes the Master

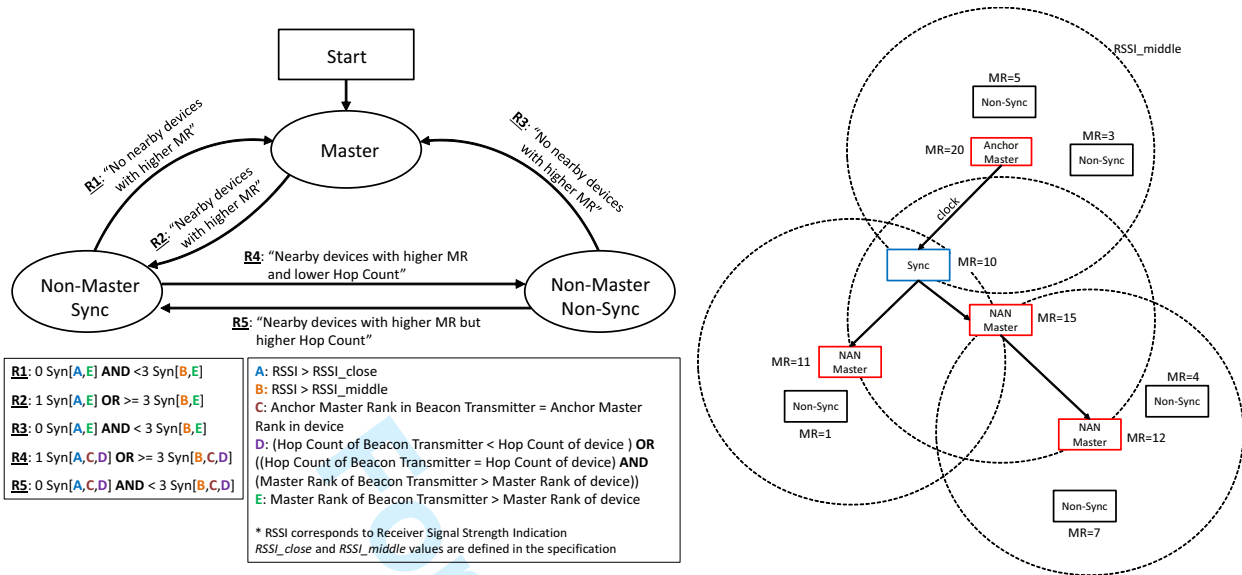


Fig. 2. The left part depicts the NAN state machine, where high level transition rules are given in the arrows between states while the formal rules R_i are described underneath. For each rule R_i the expression $n \text{ Syn}[A, B, C]$ reads as *reception of n Synch. Beacons, where each Beacon complies with conditions A, B and C*. The right part depicts an example of NAN state allocation given a specific topology and Master Rank (MR) values.

role in case it does not detect any Master device nearby. Transitions between Non-Master Sync and Non-Sync states depend mainly on the Hop Count to Anchor Master (a measure of the distance between each device and the Anchor Master); that is, the device with lower number of hops to the Anchor Master is more likely to be chosen to propagate synchronization information. A Non-Master Non-Sync device changes to Sync role in case it does not detect any Non-Master Sync device nearby. Notice that, although the term *nearby* is used loosely in the text, detection of Master and Non-Master Sync devices is actually based on objectively measurable data: the signal strength of received Synchronization Beacon frames; if the signal strength exceeds a given threshold, the transmitter of that Beacon is considered to be close. The left part of Figure 2 illustrates the NAN state machine and the exact transition rules between the different states, and the right part of Figure 2 illustrates an example NAN state allocation given a topology and Master Rank allocation. The interested reader is referred to [1] for a detailed description of the NAN state transition rules.

C. NAN Synchronization

NAN synchronization comprises the mechanisms that enable all devices in a cluster to find the time and channel (i.e. the DW) on which they should meet to announce or to discover available services. Synchronization is designed to maintain DWs among devices aligned, in order to achieve reduced discovery latency with minimum power consumption and medium occupancy.

NAN synchronization can be understood as a procedure whereby the clock reference of the Anchor Master is propagated throughout the NAN cluster by means of a subset of selected nodes (Master and Non-Master Sync devices) that form a tree structure rooted at the Anchor

Master. Note in the right part of Figure 2 how such clock distribution occurs among devices in Master state.

Synchronization relies on the transmission and processing of Synchronization Beacon frames, sent by the Anchor Master, Master and Non-Master Sync devices. The Synchronization Beacon frame is based on the original IEEE 802.11 Beacon management frame format, limited to 128Bytes, and modified to include different NAN attributes, namely the device's Master Preference, Cluster ID, Anchor Master Rank, number of hops to the Anchor Master and Anchor Master Beacon Transmission Time (AMBTT). As explained in III-E, given the key role of Synchronization Beacons, they are given a higher transmission priority.

The information contained in Synchronization Beacons is used to determine the Anchor Master and, hence, the time reference to which all NAN devices in the same cluster must synchronize their clocks. First, in order to ease the convergence of the Anchor Master selection algorithm, all NAN devices keep a record of the current and previous Anchor Masters. The latter is kept to detect stale Synchronization Beacons referencing an older Anchor Master. The current Anchor Master Record includes the rank of the current Anchor Master, the Hop Count, and the latest observed AMBTT. Then, the selection of the Anchor Master is as follows.

Any NAN device will adopt a new Anchor Master (and its clock reference) upon reception of a Synchronization Beacon announcing an Anchor Master in the same cluster whose rank is higher than that of the current Anchor Master. On the other hand, Beacons referencing other Anchor Masters with lower rank than the current one are disregarded. Note that, since NAN devices change their Master Rank every 1 to 2 minutes by modifying the Random Factor, the current Anchor Master may eventually show a lower rank in Synchronization Beacons and hence it may lose its Anchor

Master status. Finally, Synchronization Beacons referencing the current Anchor Master are used to maintain clock synchronization, but also to update the Hop Count, which reflects the position of a device within the synchronization tree, and the AMBTT value, which represents the most recent Synchronization Beacon that has been observed in the synchronization tree.

NAN devices adjust their internal clocks using the timestamp present in the appropriate Synchronization Beacons, that is, Beacons that: 1) belong to the same cluster, 2) are not stale (AMBTT values are used to state the freshness of the Anchor Master's Synchronization Beacons), and 3) reference the highest-ranked Anchor Master. However, when none of the received Beacons in a DW meet all these three conditions, NAN devices employ a default rule to adjust their clocks to the highest TSF present among the Synchronization Beacons received during the last DW, belonging to the same cluster and referencing the current Anchor Master.

NAN synchronization also considers the case when the Anchor Master becomes missing (e.g. when the device is switched off or moves away). Under that circumstance, none of the Beacons sent in the cluster will contain new (i.e. larger) AMBTT values. After three consecutive DWs without updating the AMBTT value of the current Anchor Master Record, a NAN device will assume itself as the new Anchor Master.

D. NAN Cluster selection and merging

Current Wi-Fi standards do not mandate scanning behavior, and let the decision of what network to join as an implementation choice. This approach, though, would jeopardize the ultimate goal of NAN that is to create ad-hoc clusters of synchronized devices, even when there is no relation of trust among them beyond the NAN protocol itself. Therefore, NAN specifies a cluster selection algorithm that ensures that devices will converge to a common cluster thus increasing the opportunities of discovering wanted services.

Upon discovering one or more already existing NAN clusters through the scanning of Discovery Beacons, a NAN device joins the cluster with the highest Cluster Grade (CG) value and adopts the corresponding cluster parameters, such as Anchor Master information and TSF. In [1] CG is computed as a function of the Master Preference of the Anchor Master, and the cluster TSF¹, which are both carried in Discovery Beacons. CG is very likely to be unique, and can therefore be used to arbitrate cluster selection.

Cluster merging is realized when a NAN device participating in a cluster discovers a new cluster with a higher CG. Then, the device leaves the current NAN cluster and joins the new one. In case the NAN device operates in Master or Non-Master Sync role in the previous cluster, it sends one Synchronization Beacon containing the information of the new cluster in the DW of the old cluster, hence triggering the merging process for the devices in the cluster

with lowest CG grade. Note that the size of a cluster can be limited by setting a maximum allowed Hop Count to Anchor Master. This threshold is implementation specific.

Hence, two or more NAN clusters will eventually² merge into a common cluster when their areas of influence overlap. The members of two overlapping clusters will converge to a single cluster, thus allowing the exchange of service information over a wider audience and, at the same time, reducing medium occupancy when two decoupled DWs turn into one.

E. NAN Operating in the Discovery Window

In a DW, Synchronization Beacon and Service Discovery frames are transmitted. Service Discovery frames, which can be transmitted by any NAN device regardless of its role, are used to announce services to other stations and to look for services offered by other devices in the cluster.

In order to achieve efficiency and scalability of frame transmissions in the DW, different transmission rules are followed depending on the type of frame to be transmitted. More precisely, the transmission of Synchronization Beacons is prioritized over Service Discovery frames due to the fact that synchronization is fundamental for the correct operation of a NAN cluster.

NAN transmissions are fully compliant with the IEEE 802.11 [4], although additional rules apply. Before initiating a frame transmission in the DW, each NAN device senses the channel during a time period called Distributed Inter-frame Space (DIFS). Afterwards, a backoff counter is set to a value uniformly chosen in the interval $[0, CW]$, where CW is called Contention Window. In case the backoff counter does not arrive to zero before the end of the current DW, frame transmission is aborted.

Therefore, prioritization of frame transmissions in the DW can be achieved through the selection of the corresponding backoff counter: a larger CW value is employed for Service Discovery frames, thus increasing the probability that Synchronization Beacons are transmitted at the beginning of the DW. Besides, a NAN device suspends the backoff counter for a Service Discovery frame whenever there is a Synchronization Beacon waiting for transmission.

In order to quickly propagate updated synchronization information within the DW, prioritization among Synchronization Beacons is also necessary (recall that Beacons are transmitted by different devices in a NAN). The highest priority is given to the source clock by assigning the shortest CW to the Anchor Master. In addition, Beacon transmission is scheduled as a function of the device's Hop Count value, thus allowing devices to first receive an updated timestamp from devices higher in the synchronization tree (i.e. closer to the Anchor Master), and then forward this updated timestamp to the lower levels of the synchronization tree.

On the other hand, concurrent transmission of Service Discovery frames within a DW may potentially result in a

¹ $CG = 2^{64}A_1 + A_2$, with A_1 the Master Preference of the Anchor Master, and A_2 the 8-octet TSF value of the NAN cluster.

²Notice that merging time cannot be guaranteed as it heavily depends on the dynamics of the involved devices.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

large number of collisions when the number of synchronized NAN devices is high. Hence, in order to mitigate intra-cluster collisions, NAN devices employ a large CW to contend for channel access ($CW = 511$). However, such a large CW may prevent NAN devices from transmitting when legacy Wi-Fi traffic using a small CW ($CW = 15$) is present. Therefore, for service discovery frames the following algorithm is used. In addition to the previous backoff: NAN devices run a second backoff counter that uses a small CW value ($CW = 15$), but it is only started if the frame could not be transmitted before a given deadline, which is randomly chosen between the beginning and the end of the DW. After this deadline, only the backoff counter having the shortest value is considered. Notice that this double backoff strategy effectively avoids intra-cluster collisions, while ensuring transmission opportunities even with legacy traffic and reducing unnecessary long backoff delays when only few NAN devices are present.

Finally, in order to increase channel efficiency, different service descriptors (i.e. service announcement or request) coming from the same NAN device can be aggregated in a single Service Discovery frame. In addition, since a DW can accommodate only a limited number of transmissions, the NAN specification presents a procedure according to which a NAN device transmits one Service Discovery frame in only a subset of the DWs, thus minimizing the number of transmitted frames in each DW.

IV. THE NAN DISCOVERY ENGINE

A. NAN Service Model / API

As illustrated in Figure 1, services and applications communicate with the NAN Discovery Engine (DE) in order to access to service information through the usage of service primitives. On the other hand, the DE communicates with the NAN MAC, which is responsible for handling NAN Service Discovery frames.

Services are identified by their Service ID, which is a 6 Byte long hash of the service name (a UTF8 string uniquely identifying the service). The duration of a Service Discovery frame is limited to $400\mu s$, which allows other devices' transmissions in the DW and thus assures network scalability.

There are two basic NAN service primitives, which are carried in NAN Service Discovery frames: Publish and Subscribe. Publish-related methods are used to make a service discoverable for other devices. A call to the Publish method can be translated either to a periodic broadcast of Publish messages announcing the service, or limited to the generation of a response only when a Subscribe message is received for that service. Subscribe methods allow NAN devices to search for a given service. The Subscribe function may be configured to operate either in passive (waiting for corresponding Publish messages sent by other devices) or in active mode (transmitting Subscribe messages).

Following the discovery of a service, a NAN device may need to establish a connection with a peer device outside the

NAN. The NAN Connection Capability attribute, present in Service Discovery frames, may assist the connection setup. In this way, the NAN Connection Capability attribute informs about the different Wi-Fi-based connection methods supported by that NAN device (e.g. through WLAN infrastructure, Wi-Fi Direct, etc.).

B. Security Aspects

There are several security aspects to take into account for the correct operation of a NAN. Security in a NAN needs to be built into the applications using the NAN primitives, and is independent of IEEE 802.11 MAC security.

The licit transmission of Synchronization Beacon frames in a NAN is fundamental. A malicious device could disrupt the synchronization process by sending corrupted Beacons including false cluster information. Due to the fact that synchronization information is broadcast by several devices in a NAN (cf. Section III-B), the process is robust enough to overcome individual and specific malicious Beacon frames. However, a continuous transmission will lead to a denial of service. Mechanisms to detect false synchronization information are implementation specific.


Discovery is an inherently open process where Publish and Subscribe methods include Service Identifiers. The Service Identifiers are truncated hashes of a more humanly readable Service Name. The Service Identifiers are opaque and if not known by a device the identifiers are not readily identified as belonging to a particular application. This opacity allows private groups to be formed that use their own unique Service Identifiers. The group unique identifier can be created by mixing a shared group key with the Service Name, thus defining a group specific Service Identifier. Confidentiality of the information carried in NAN frames may be supported by applications. Encryption of NAN fields and the required group key distribution are out-of-scope of the NAN specification and are managed by the applications using NAN.

Privacy is an important security concern affecting NAN devices. The fixed MAC address of a Wi-Fi device enables easy identification and tracking of users. To improve Wi-Fi privacy, the NAN specification allows the use of randomly selected local MAC addresses, which should be changed occasionally to avoid address tracking. The means and frequency of address changing are implementation specific. The occasional changing of MAC addresses should not interfere with the inherent robustness of the NAN synchronization procedure.

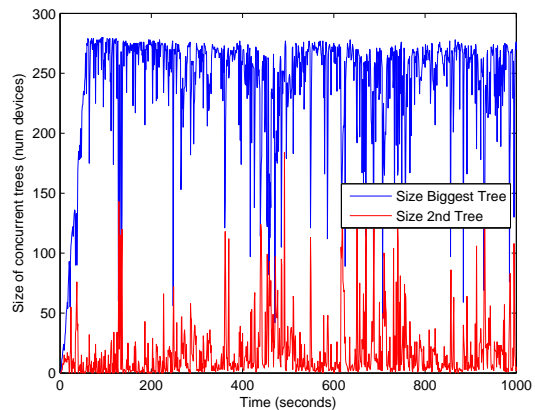
V. PERFORMANCE EVALUATION

A. Simulation Set up

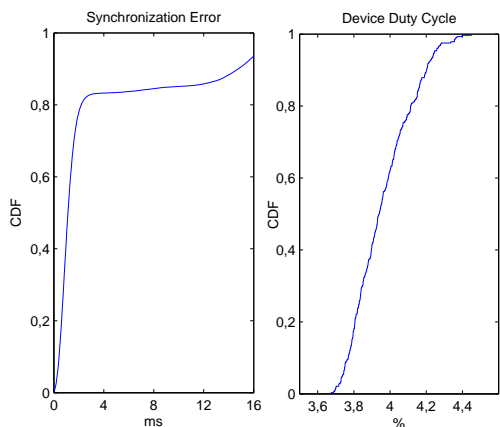
In this section we use packet level simulations to illustrate the performance of the Wi-Fi NAN technology in a realistic scenario. In particular, we consider a set of 300 users carrying a Wi-Fi NAN enabled device moving at pedestrian speeds through the streets of Osaka downtown. Our packet level simulator is based on OPNET [5], where we have developed the protocols described in Section III,




(a) Example of NAN synchronization tree in Osaka downtown



(b) Size of concurrent synchronization trees over time



(c) CDF of synchronization errors and device duty cycle



(d) CCDF of time (%) spent by a NAN device in each NAN state

Fig. 3. NAN performance in Osaka downtown

and on MobiReal [6], which is used to generate realistic mobility patterns. Each simulation run represents 3000 seconds of simulated time, and multiple runs are considered to ensure that results are statistically significant.

In our simulations each device randomly selects a Master Preference value that is updated every T , with T chosen randomly by each device between 2 and 10 minutes. We consider a clock drift of ± 500 ppm, and each NAN device performs passive scanning for 200 milliseconds with an interval randomly chosen between 10 and 20 seconds. In addition, the transmission power is fixed to 32 mW and the receive sensitivity to -84 dBm.

B. Results

We start discussing the performance of Wi-Fi NAN by looking at Figure 3(a) that depicts a snapshot of a typical NAN synchronization tree obtained by means of the algorithm described in section III-C, where a tree is a set of NAN devices following the same Anchor Master device. Figure 3(a) depicts, for a particular time instant, the position of each user in Osaka downtown as dots, and the relation between a NAN device and its parent in the synchronization tree as solid lines, where the parent device is the device

that sent the Beacon used for synchronization. As a result of our experiment we observed that in the considered scenario the maximum number of hops between any device and the Anchor Master was found to be five. In addition, we observed that devices tend to cluster around a single synchronization parent, which is a device in Master state. Recall from section III-B that within a local neighborhood only the device with the highest Master Rank transitions to Master state. It is also worth noting in Figure 3(a) that some devices (dark red dots) are isolated. These are devices that lost track of the Anchor Master and reset themselves in Anchor Master state.

In order to further understand the dynamics of the synchronization trees formed in NAN, Figure 3(b) depicts over time (limited to 1000 seconds for clarity) the size of the two biggest synchronization trees that exist concurrently in our scenario. We can see in Figure 3(b) how most devices tend to follow the same Anchor Master and belong to the same tree (blue line), which is the goal of the NAN protocol. However, a small percentage of devices, while belonging to the same cluster, temporarily follow a different Anchor Master. This behavior is to be expected in an extremely dynamic environment like the one considered

in our experiments where devices may lose sight of their Anchor Master Beacons, due to mobility or because of updates in the Master Rank values.

The left part of Figure 3(c) illustrates, with a cumulative distribution function (CDF), the time difference between the wake-up times of each device in our scenario for all DWs across several simulation runs. Notice that, if all devices would wake up at exactly the same time, this CDF would be zero. However, the previous is not possible given that a clock tolerance of ± 500 ppm may introduce a clock drift of up to 0.5 milliseconds between DWs. Indeed, we see in Figure 3(c) that 80% of the devices have a synchronization error below 2 milliseconds, which guarantees a correct operation in the DW. However, the remaining 20% of devices experience larger synchronization errors due to the fact that they may be temporarily following a different Anchor Master, as illustrated in Figure 3(b). In addition, we measured how much time of a DW is spent in the transmission of the Synchronization Beacons required to maintain synchronization in the NAN cluster. Our results showed that in 90% of the DWs the Synchronization Beacons overhead is below a 14% which, given the previous synchronization performance, guarantees that there is enough effective time within a DW for NAN devices to exchange Service Discovery frames.

The right part of Figure 3(c) illustrates, with a CDF, the duty cycle experienced by the NAN devices in our experiment, where we can observe a median duty cycle around 4%, which is a worst case in practice because, in our experiment, all devices (even those in Non-Master Non-Sync state) wake up and listen every DW. Considering a modern mobile device battery capacity of 2600 mAh [7] powered at 4 V [8], and the Wi-Fi chipset power consumption model used in [3], the duty cycle CDF curve in Figure 3(c) can be translated to a power consumption CDF curve, which we do not include for the sake of space, from which it can be derived that, disregarding other sources of power consumption, devices in our experiment could afford more than ten days of continuous NAN operation, hence validating the original intent of the NAN protocol of achieving a continuous background operation.

Finally, Figure 3(d) illustrates with a complementary CDF the probability that a NAN device spends a certain percentage of its time in a given NAN state. Recall that the NAN protocol described in section II consisted of four different NAN states: Anchor Master, Master, Non-Master Sync and Non-Master Non-Sync. The purpose of the NAN states is to let devices share in a fair way the burden of generating Beacons, while allocating a higher share of work to devices with higher Master Preference values. Consequently, we can see in Figure 3(d) how devices tend to operate most of their time in Non-Master Non-Sync state, which validates the design of the NAN protocol.

operation of these radios could be made energy efficient, these devices could continuously advertise and discover interesting services on behalf of their users. Wi-Fi Neighbor Awareness Networking (NAN) is a novel technology being developed in the Wi-Fi Alliance that attempts to solve this problem. In this paper we have provided a thorough overview of the MAC layer mechanisms that underpin this technology, and have provided a performance evaluation that illustrates the performance to be expected from future NAN devices in realistic scenarios.

REFERENCES

- [1] Wi-Fi Alliance, *Wi-Fi NAN Technical Specification*, TG Baseline r11, January 2014.
- [2] W. Ye, J. Heidemann and D. Estrin, *Medium Access Control with Coordinated, Adaptive Sleeping for Wireless Sensor Networks*, IEEE/ACM Transactions of Networking, vol. 12, no.3, pp. 493-506, 2004.
- [3] Camps-Mur, D.; Loureiro, P., *E²D Wi-Fi: A Mechanism to Achieve Energy Efficient Discovery in Wi-Fi*, Mobile Computing, IEEE Transactions on , vol.PP, no.99, pp.1,1. 2014.
- [4] IEEE 802.11 Task Group, *IEEE Standard for Information technology Specific requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, 2012.
- [5] OPNET, <http://www.opnet.com>.
- [6] K. Maeda, K. Sato, K. Konishi, A. Yamasaki, A. Uchiyama, H. Yamaguchi, K. Yasumoto and T. Higashino, *Getting Urban Pedestrian Flow from Simple Observation: Realistic Mobility Generation in Wireless Network Simulation*, in Proceedings of the 8th ACM/IEEE International Symposium on Modeling, Analysis and Simulation of Wireless and Mobile Systems, MSWiM, 2005.
- [7] Samsung Galaxy S4, <http://www.samsung.com/us/mobile/cell-phones-accessories/EB-B600BUBESTA>
- [8] Falaki, Hossein, et al. *Diversity in smartphone usage*, Proceedings of the 8th international conference on Mobile systems, applications, and services. ACM, 2010.

VI. CONCLUSIONS

There is an untapped potential in the Wi-Fi radios embedded in smartphone and tablet devices. If background