Content Dissemination in Wireless Networks Exploiting Relaying and Information-Centric Architectures

Pantelis A. Frangoudis*, George C. Polyzos§, and Gerardo Rubino*
*INRIA Rennes-Bretagne Atlantique
§Mobile Multimedia Laboratory, Department of Informatics, Athens University of Economics and Business
Email: pantelis.frangoudis@inria.fr, polyzos@aueb.gr, gerardo.rubino@inria.fr

***Abstract*—We focus on the problem of efficiently integrating wireless users in future Information-Centric Networks (ICN), where communication is based on publish-subscribe primitives. The current host-centric Internet paradigm is abandoned in favor of information-oriented, rendezvous-based communication, where multicast data delivery is the norm. However, Wi-Fi, the predominant means of local wireless connectivity today, but also 3G and 4G technologies, are known to suffer from poor multicast performance. Data destined to a broadcast or multicast address are typically transmitted at lower rates to increase reliability for clients with poor signal conditions, causing unfavorable delays for high-rate users. One approach to this problem is to designate a subset of the clients as relays who re-broadcast packets for other clients at a higher rate. Given that different types of content have different performance requirements, we exploit content-awareness, inherent in our environment, to optimize for different criteria on a per-content basis. For this purpose, we provide a multi-objective optimization formulation for the problem of relay selection and rate assignment, which can capture the tradeoff among reliability, performance and energy cost.**

I. Introduction

In recent years, we have pushed towards "clean-slate" *information-centric* Internet architectures. The current host centric model is abandoned for an information-centric one, where named data, instead of named hosts are the core of the new communication paradigm.

This shift was motivated by the observation that the Internet architecture has not changed fundamentally since its inception, but user (stakeholder) behavior and application demands have all dramatically changed. Much of the traffic nowadays involves content dissemination via CDNs or proxies, which mediate communication between content publishers and consumers, placing focus on the information itself, rather than on the communication endpoints. The same view is manifested by the popularity of applications such as Bittorrent, where one is not interested in *who* provides content, but in *what* is exchanged. This way of thinking has given rise to clean-slate internet designs which are based on information-centric principles.

At the same time, with the low cost and high speed of wireless technologies for the home network and the proliferation of mobile devices with wireless networking capabilities, much of this information is delivered to users over wireless links. In this work, we consider wireless content delivery over an information-centric internetworking infrastructure, and, in particular, explore ways to improve performance by means of *relaying*.

Our work is put in the context of the Publish-Subscribe Internet (PSI) architecture [1]. PSI approaches information-centrism by applying the *publish/subscribe* principle at all networking layers. The main PSI entities are *publishers* and *subscribers* and their communication is brokered by special *rendezvous* nodes. PSI and ICN in general has various advantages and features such as built-in support for caching and mobility [2], multihoming, multipath, and security [3].

The architecture makes few assumptions about the lower communication layers. In particular, a basic expectation is that at the PHY/MAC layer there is a broadcast protocol. An example scenario involves subscribers receiving publications of content availability and broadcasting their intent to receive it. Also, much of the data (content) traffic is expected to be of broadcast nature and the architecture can promote and enforce this.

We focus on a particular networking scenario where subscribers are attached to Wi-Fi Access Points (APs). (However, the problems and the suggested approach of relaying is general, applies to and has been suggested for most other wireless technologies such as UMTS/3G, 4G/LTE-LTE Advanced etc.) Multiple users within a Wi-Fi cell can subscribe for the same content, which is eventually multicast to them. We incorporate APs in the PSI architecture to make them be aware of the traffic to and through them and so that they can distinguish between publications, subscriptions and publication data.

Local broadcast and multicast, typical in a PSI world, work well with a reliable Ethernet substrate. However, things are quite different when publishers or subscribers are attached to a Wi-Fi link. The fundamental problem is that IEEE 802.11 was not designed for scenarios where broadcast traffic dominates [4]; the backoff window is never increased, since there are no acknowledgements for broadcast packets, and typically the transmission rate is lower compared to unicast packets to achieve more reliable delivery for low-rate users.

In our target environment, multiple clients attached to an AP subscribe for multiple content publications outside the Basic Service Set (BSS). Publication data would be broadcast by the AP and subscribers would receive them, with non-interested

QSHINE 2014, August 18-20, RHODES, Greece

DOI 10.4108/icst.qshine.2014.256342

clients simply ignoring them. Packets would be transmitted by default at a low rate, irrespective of their content and their subscribers. Instead, we propose that with awareness of each packet's content and intended recipients, we can smartly pick specific high-rate clients to act as relays for other co-subscribers whose connection with the AP is of worse quality. Thus, instead of a single broadcast transmission at a low rate, multiple broadcast transmissions of the same content at a high rate will take place. If relay selection is careful, both throughput and reliability advantages are possible.

A question that naturally emerges is whether content-awareness can help us in achieving a more efficient (or more flexible) relaying scheme. State-of-the-art relaying mechanisms for Wi-Fi networks [5], [6] select relays solely based on rate capabilities. Considering broadcast traffic, a selected relay node would always re-broadcast packets to its neighbors, whether or not these packets are of interest to them. We argue that by exploiting content-awareness, we can achieve better performance by building per-content relay schedules and even save energy by selectively putting clients in power-save mode when appropriate (e.g., when data they have not subscribed for are multicasted)[1].

To address the potentially conflicting requirements that different types of content have, we provide a multi-objective optimization framework which can be used by relay and rate selection algorithms to optimize for reliability, performance or energy consumption on a per-content basis.

This paper is structured as follows: In Section II we provide a review of related works and background information on the PSI architecture. We present our design towards relay-based data delivery in a wireless ICN environment in Section III, and propose a multi-objective optimization framework for relay selection and transmission rate assignment for this design in Section IV. We discuss ongoing and future work and conclude the paper in Section V.

II. Background

A. State of the art

Since performance anomalies emerge in IEEE 802.11 networks when there are stations which use low rates [7], solutions that opportunistically exploit specific clients as repeaters to reduce low-rate transmissions have been proposed. Bahl et al. [5] design and implement *SoftRepeater*, a Wi-Fi-compatible software solution where high-rate stations behave as repeaters for low-rate ones when it is beneficial to do so.

A system which addresses multicast traffic and also focuses on smart relay selection is PeerCast [6]. PeerCast handles multicast packets in batches; the first packets of each batch are sent at varying rates and are simultaneously acknowledged by clients upon reception. Since low-rate stations are not expected to receive packets sent at high rates, the amount of acknowledgements (and, thus, the received power detected at the AP) increases as the transmission rate decreases. After each batch, clients serially send a report on the packets successfully received and these reports are overheard by other stations. These mechanisms help the AP decide which clients should act as relays. It therefore appears that PeerCast selects relays based on rate criteria. On the contrary, our design aims to explore alternative relay selection objectives and, in particular, put information-centrism and content-awareness at the center: Relay plans can be built per publication and different optimization criteria can be applied taking information semantics into consideration.

We focus on the integration of wireless users in PSI, a clean-slate, purely information-centric architecture for the Future Internet based on the application of the publish-subscribe primitives throughout the "networking stack." In this information-centric networking scenario, multicast is the norm. Since delivery of such traffic can be problematic in modern wireless networks (with physical layers that optimize and exploit point-to-point delivery through various means, including directionality, array antennas, MIMO etc.), we study how to exploit content-awareness to facilitate efficient multicasting in a setting where content subscribers connect to the pub/sub Internet over wireless links. A brief description of the PSI architecture and its principles are given in Section II-B. More details are available in [1], [2], [8], and the deliverables of the PSIRP [9] and PURSUIT [10] projects.

Other projects with similar motivation also exist. CCN [11]/NDN [12] is a research effort that also aims to design an information-centric Internet. In CCN/NDN, routing is based on hierarchical naming. Consumers broadcast *interest* packets that contain the name of the content in request and *data* packets whose content name is a suffix of the name in the interest packet are assumed to satisfy this interest. 4WARD [13] and SAIL [14], other related ICN EU-funded projects, aim at allowing various types of networks to coexist and cooperate in a smooth and cost-efficient manner.

B. The PSI architecture

1) Information identification and scoping: In PSI, information items are identified by statistically unique labels, which are used to match subscriptions with publications (*rendezvous* function). Such labels are called Rendezvous Identifiers (RId). Identifiers are flat and endpoint-independent, since, in an information-centric networking environment, the basic design premise is that location is decoupled from identity.

Information is organized in a hierarchy of *scopes* and any information item is also identified by a Scope Identifier (SId), a flat and endpoint-independent label (as is the case for RIds). A scope is a classification of information in groups with similar semantics, visibility properties, etc. Information is always published under (at least) one scope, which is set by the publisher.

2) Core functions: PSI design follows a recursive approach. Each layer implements a set of core functions, utilizing the same core functions of the lower layers. These functions are

[1]A similar but less pronounced effect can be achieved in an IP world by having APs understand and participate in IGMP (e.g., implement IGMP router functionality). We consider here a non-IP world and focus on how to exploit information identification/data naming (which is a key property/global requirement in ICN), which brings these features at the heart of the architecture.

Rendezvous, *Topology management* and *Forwarding* (RTF). Network composition is performed recursively: RTF is implemented per scope, first locally, then at the LAN level, at the WAN level, and so forth.

The Rendezvous function is responsible for matching publications with subscriptions. The node where the matching of a publisher's content with a subscriber's interest takes place is referred to as the rendezvous point (RP). RPs initiate routing, forwarding, and distribution decisions, eventually leading to the delivery of the content from publishers to subscribers.

The Topology function monitors the network topology, detects changes and is responsible for creating information delivery paths. Upon a successful publication-subscription match, the rendezvous component can request forwarding information from this module. Policies and specific dissemination strategies can be put in effect by the Topology function to build delivery structures.

The Forwarding function implements information forwarding through the paths dictated by the Topology modules. Forwarding is implemented using the LIPSIN approach [15]: Each link is assigned a *forwarding identifier (FId)* and the topology module builds a delivery tree which includes such FIds. FIds are encoded in a Bloom filter (*zFilter*, in the LIPSIN terminology) which is placed in the header of each individual packet. The forwarding module of each node, upon receiving a packet, ORs each of its outgoing link IDs with the zFilter to decide where to forward the packet to. This approach naturally supports multicast.

III. Relay-based data delivery

A. Requirements

In the scheme that we propose, there are two basic requirements for each AP:

- Awareness of the content each of its clients has subscribed for.
- Knowledge about the rates at which clients are capable of communicating with the AP and with each other.

With the above information in place, the AP can come up with an efficient relaying schedule so that low rate transmissions within its cell are minimized. We show how the above information can be collected, managed and applied for efficiently planning relay transmissions.

B. Construction of a rate graph

The AP is responsible for maintaining a BSS-wide *Rate Graph (RG)* which encodes the rate capabilities between pairs of nodes. We define the rate graph $RG = (V, E)$ as the set of vertices $V = v_1, v_2, ..., v_n$ which correspond to BSS nodes (AP and clients) and edges E between these vertices. An edge $(v_i, v_j) \in E$ iff nodes v_i and v_j can communicate at any of the available set of rates. The weight of this edge is the maximum rate at which the two nodes can communicate. An example WLAN topology and the corresponding RG are shown in Fig. 1. Stations B and C are "close" to the AP and the latter can communicate with them at a high rate (e.g., 54Mbps for IEEE 802.11g), while a lower rate should be used to transmit packets to nodes D and E. Still, there are clients which are capable of transmitting packets to D and E at a high rate (C and B respectively) and could be exploited to relay packets to them on behalf of the AP.

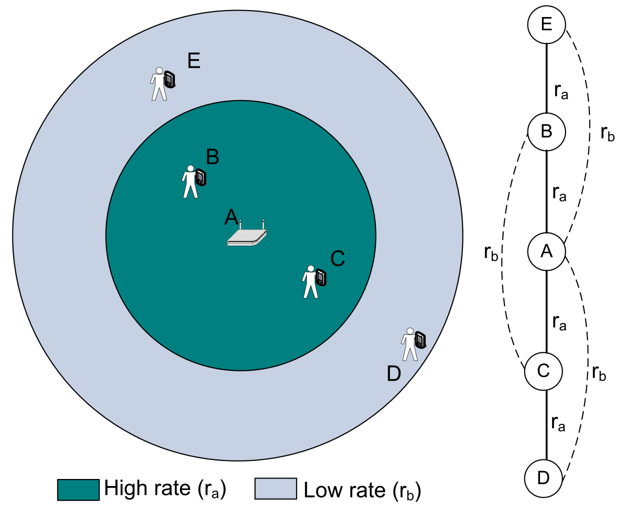

Fig. 1. Example WLAN topology and the respective rate graph. For presentation clarity we do not show the communication ranges of clients.

We assume for the moment that subscription traffic is upstream and publication and data traffic is downstream. Clients periodically monitor the channel and when they detect a subscription message, they record the signal power of the subscription and the identity (e.g., MAC address) of the sender. At the same time, they overhear downstream messages from the AP (beacon frames, publications, publication data). After the monitoring period ends, each client submits $< ID, SNR >$ tuples to the AP at a low rate.

The AP adds edges to the RG based on the reported SNR values, which are mapped to achievable transmission rates. This implies that our rate graph construction scheme is approximate. Measurement studies [16], [17] have shown that physical layer metrics are not always good estimators of delivery probability and, in turn, achievable rates. However, correlation exists and such an approach serves reasonably as a starting point [5].

Also, it should be noted that we assume a certain degree of symmetry; if we estimate that node B can transmit to A at rate r_i based on the SNR of packets intercepted by node A, we assume that the same rate can be achieved for $B \rightarrow A$ transmissions.

C. Maintenance of per-client subscription state

The AP keeps state regarding the content each station has subscribed for. This information will be used when coming up with a transmission/relay plan. Recall that, in the scenario we study, the AP is the first to receive content subscriptions before they are further propagated to the information-centric network. Conversely, it broadcasts publications and delivers content within its BSS. By inspecting this traffic, it maintains a $m \times n$ table, where m is the number of stations associated with it and n the number of publications that it has broadcast to its Wi-Fi cell. A table cell is marked if the respective client has subscribed for the respective publication.

D. Selecting relays

The Rate Graph and the subscription table are the input to a relay scheduling algorithm for multicasting content to interested stations. This algorithm is executed per publication. Namely, a different relay plan is picked based on the set of interested parties per information item. After selecting a number of relays and the respective transmission rates, the sequence of transmissions is decided and the schedule is broadcast by the AP at a low rate. When a client receives the relay schedule, it maintains state about the rate at which it is required to re-broadcast the data it has received. Since relay selection is per publication, a node may at the same time serve as a relay for a subset of the publications it has subscribed for at different rates each. The relay selection algorithm can be executed periodically or when changes are detected in the rate graph or the subscription table.

E. Managing content requirements

Based on the nature of the requested content, different objectives should be attained. These objectives can be conflicting. For example, reliable delivery may require low-rate transmissions to maximize the number of nodes receiving the content, while at the same time increasing delivery time for specific users who are reachable at higher rates. Also, specific constraints could be in place, such as minimum coverage requirements or energy constraints for battery-powered devices. The PSI architecture facilitates managing this information using the scoping mechanism: Specific scopes can be defined, each encoding the relative importance of reliability, delivery time and energy cost for the content items belonging to it, as well as the respective constraints.

IV. An optimization framework for relay selection and rate assignment

A. Problem formulation

To cater for conflicting requirements for information delivery, we propose a framework for relay selection and rate assignment which can optimize for different criteria on a case-by-case basis, applying a multi-objective optimization problem formulation. We introduce the following notation:

- $R = \{r_1, r_2, \ldots\}$ is the set of available rates. $r_1 = 0$ denotes no transmission.
- A is the set of the k potential transmitters (AP and potential relays).
- $e(i)$ is the energy cost for the transmission of a content item for node i.
- $S_{i,r}$ is the set of nodes reachable by node i when transmitting at rate r.

The purpose of our scheme is to select a transmission plan $p = [a_1 \ldots a_k]$, where $a_i \in R$ is the assigned rate for potential transmitter i, considering the following optimization criteria for data delivery.

Completion time: This is an expression of the time it takes to deliver 1 bit under a specific transmission plan. With the assumption that transmissions are carried out serially (i.e., no simultaneous transmissions are possible), the completion time under transmission plan p is given by:

$$T(p) = \sum_{\substack{i=1 \\ a_i > 0}}^{k} \frac{1}{a_i}. \tag{1}$$

There are cases where we can sacrifice reliable delivery in the interest of speed. Namely, we may be more interested in serving a subset of the clients at the highest possible rate, even if this would mean that others could suffer data losses. In this case, low completion time should have more weight in the relay selection process.

Reliability: Conversely, there are cases when reliable delivery is the main objective. Our expression of reliability is the number of clients capable of receiving the transmitted information under a specific transmission plan:

$$C(p) = \left| \bigcup_{i=1}^{k} S_{i,a_i} \right|. \tag{2}$$

Energy cost: In typical WLAN settings, some of the clients are on battery power. Selecting these devices as relays means that their battery will be exhausted sooner. Thus, minimizing the total energy cost, given by the following formula, is another objective:

$$E(p) = \sum_{i=1}^{k} H(a_i) e(i), \tag{3}$$

where $H(x) = \begin{cases} 1 & \text{if } x > 0 \\ 0 & \text{otherwise} \end{cases}$. We use a simple model where the energy cost for the transmission of a content item is constant, irrespective of the rate or size of the item, therefore

$$e(i) = \begin{cases} c & \text{if node } i \text{ is on battery power} \\ 0 & \text{otherwise} \end{cases}. \tag{4}$$

Ideally, the system designer should aim to maximize reliability, while minimizing completion time and energy cost. At the same time, specific constraints may need to be satisfied. For example, a minimum number of clients should receive publications under a specific scope, while for another scope it is possible that a deadline for the delivery of requested information algorithm should be respected even at the expense of not covering some clients (e.g., in the case of a real-time multimedia streaming application). Also, energy constraints could be defined. We formulate this optimization problem as follows:

$$\begin{aligned} \underset{p}{\text{maximize}} \quad & F(p) = \begin{bmatrix} C(p) \\ -T(p) \\ -E(p) \end{bmatrix} \\ \text{subject to} \quad & C(p) - C_{min} \geq 0, \\ & T_{max} - T(p) \geq 0, \\ & E_{max} - E(p) \geq 0. \end{aligned} \tag{5}$$

C_{min} denotes the minimum size of the set of covered clients, T_{max} is the maximum allowed completion time, and E_{max} is the maximum allowed energy cost.

We define vector $W = [w_R\ w_T\ w_E]$, where w_R, w_T, and w_E are the weights of the reliability, completion time and energy cost objectives and $w_C + w_T + w_E = 1$.

Since there is typically no solution which optimizes all objectives simultaneously, a *scalarization* [18] approach can be applied, where the multiple components of the objective function are appropriately weighted and combined to a single objective function, and a scalar optimization problem is solved instead. The weight vector expresses the importance of each criterion in the selection of the final solution among a set of Pareto optimal[2] solution vectors for the original problem. Miettinen and Mäkelä [18] provide a thorough overview of scalarization approaches.

A further issue is that the components of the objective function have different units and orders of magnitude. The functions thus need to be properly transformed via a normalization process. Marler and Arora discuss and compare various such function transformation methods [19].

B. Finding an optimal solution

Problems related to optimally selecting a set of relays to rebroadcast messages have been shown to be NP-hard by reduction to the minimum set covering problem [6], [20].

Given a specific weighting of the optimization criteria, exhaustively searching the whole solution space and finding the assignment that optimizes the objective function is prohibitively expensive computationally, even for small problem instances (e.g., for cells with few clients). Therefore, appropriate heuristics need to be put in effect. While the brute force approach has a time complexity of $O(|R|^n)$, where n is the number of WLAN nodes (clients and AP) and $|R| = 9$ for IEEE 802.11g, in practical scenarios (i) the set of available rates can be reduced (e.g., by allowing only a specific subset of transmission rates, such as $\{0, 6, 54\}$ Mbps), and (ii) the size of the set of potential transmitters k could be fixed to a small constant (e.g., allowing at most two relays). In this case, the number of potential rate assignments to evaluate is constant, and the complexity of the algorithm depends on the complexity of evaluating the objective functions. (In our formulation, this is polynomial, as the most costly operation is the calculation of the set of covered users.) Experiments with a relay selection algorithm based on our optimization framework (not presented here due to lack of space) indicate that these heuristics have both reasonable running time when executed on top of typical off-the-shelf wireless equipment, and derive solutions with improved performance compared to default IEEE 802.11g multicast schemes.

V. Conclusion

In view of a future wireless information-centric networking environment, we presented mechanisms that can fit well within the context of the Publish-Subscribe Internet (PSI) architecture, exploiting content-awareness for smart relay selection in the wireless domain. Since different types of content may have different requirements, we proposed a multi-objective optimization framework for selecting optimal relaying strategies, capturing the tradeoff between reliability, performance and energy cost. Our ongoing and future work focuses on designing relay selection algorithms with tolerable complexity for realistic WLAN sizes, and on tackling the challenges of implementing and integrating them with the PSI architecture and prototype [8]. We further study practical use cases (such as the dissemination of layered video content) which will highlight the flexibility of our approach and will demonstrate quantitatively the improvements it can bring about.

[2] A solution vector for the original problem is Pareto optimal *iff* it is not possible to move from that point and improve at least one objective function without negatively affecting any other objective function.

References

[1] N. Fotiou, D. Trossen, and G.C. Polyzos, "Illustrating a publish-subscribe Internet architecture," *Telecommunication Systems*, vol. 51, no. 4, pp. 233–245, 2012.

[2] G. Xylomenos, X. Vasilakos, C. Tsilopoulos, V.A. Siris, and G.C. Polyzos, "Caching and mobility support in a publish-subscribe internet architecture," *IEEE Communications Magazine*, vol. 50, no. 7, pp. 52–58, 2012.

[3] N. Fotiou, G. Marias, and G.C. Polyzos, "Publish-Subscribe Internetworking Security Aspects," in *Trustworthy Internet*. Springer, N. Blefari-Melazzi, G. Bianchi, L. Salgarelli, eds., May 2011.

[4] R. Oliveira, L. Bernardo, and P. Pinto, "The influence of broadcast traffic on IEEE 802.11 DCF networks," *Computer Communications*, vol. 32, no. 2, pp. 439–452, 2009.

[5] P. Bahl, R. Chandra, P.P.C. Lee, V. Misra, J. Padhye, D. Rubenstein, and Y. Yu, "Opportunistic use of client repeaters to improve performance of WLANs," *IEEE/ACM Transactions on Networking*, vol. 17, no. 4, pp. 1160–1171, 2009.

[6] J. Xiong and R.R. Choudhury, "Peercast: Improving link layer multicast through cooperative relaying," in *Proc. IEEE INFOCOM*, April 2011, pp. 2939–2947.

[7] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda, "Performance anomaly of 802.11b," in *Proc. IEEE INFOCOM*, March 2003, pp. 836–843.

[8] D. Trossen and G. Parisis, "Designing and realizing an information-centric Internet," *IEEE Commun. Mag.*, vol. 50, no. 7, pp. 60–67, 2012.

[9] "PSIRP: Publish-Subscribe Internet Routing Paradigm," http://www.psirp.org/.

[10] "PURSUIT: Publish-Subscribe Internet Technology," http://www.fp7-pursuit.eu.

[11] "CCNx," http://www.ccnx.org.

[12] "NDN," http://www.named-data.net/.

[13] "4WARD," http://www.4ward-project.eu/.

[14] "SAIL," http://www.sail-project.eu/.

[15] P. Jokela, A. Zahemszky, C.E. Rothenberg, S. Arianfar, and P. Nikander, "LIPSIN: line speed publish/subscribe inter-networking," in *Proc. ACM SIGCOMM*, August 2009, pp. 195–206.

[16] S.H.Y. Wong, H. Yang, S. Lu, and V. Bharghavan, "Robust rate adaptation for 802.11 wireless networks," in *Proc. ACM MobiCom*, 2006, pp. 146–157.

[17] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris, "Link-level measurements from an 802.11b mesh network," in *Proc. ACM SIGCOMM*, 2004, pp. 121–132.

[18] K. Miettinen and M.M. Mäkelä, "On scalarizing functions in multiobjective optimization," *OR Spectrum*, vol. 24, no. 2, pp. 193–213, 2002.

[19] R.T. Marler and J.S. Arora, "Function-transformation methods for multi-objective optimization," *Engineering Optimization*, vol. 37, no. 6, pp. 551–570, 2005.

[20] B. Mans and N. Shrestha, "Performance evaluation of approximation algorithms for multipoint relay selection," in *Proc. Med-Hoc-Net '04*, June 2004.