

Challenges and New Approaches for Efficient Data Gathering and Dissemination in Pervasive Wireless Networks

Elena Fasolo[†], Christian Prehofer^{*}, Michele Rossi[†], Qing Wei^{*},
Jörg Widmer^{*}, Andrea Zanella[†] and Michele Zorzi[†]

[†]DEI, University of Padova, via gradenigo 6/A – 35131, Padova, Italy

^{*}DoCoMo Euro-Labs, Landsberger Strasse 312 – 80687 Munich, Germany

Abstract—Motivated by the current trends in wireless technologies, we present challenging scenarios consisting of moving and uncoordinated wireless objects as well as algorithms that are likely to perform well in such environments. Our primary interest is to design network protocols to efficiently build and maintain distributed communication systems based on opportunistic design principles. Such systems should continue to operate well even if communication is sporadic, no end-to-end routes are available, and sources and consumers of information are not known in advance. Also the environment poses many challenges, that are typical for wireless mobile devices, such as wireless channel impairments, the impossibility of having up-to-date and accurate views of the network topology at every device, and limited energy resources. In the second part of the paper, we discuss and propose solutions for data dissemination in such an environment by exploiting network coding techniques. Furthermore, we report first results which reveal that network coding largely outperforms standard solutions which are based on the classical store and forward paradigm.

I. INTRODUCTION

The focus of the present paper is on a wireless network scenario where communicating elements of various nature can interact by exchanging and processing information in a distributed manner. In this environment users can both access very diverse types of information in a seamless and distributed manner and actively contribute to the creation and dissemination of content in the network. In the most general case, such a distributed environment will include the presence of sensing units, actuators, databases, and end-users which all cooperate to distribute, store, and utilize information.

The challenges that have to be faced to realize such a vision are manifold. One of the main issues is given by the inherent heterogeneity, where intelligent devices (cellular phones, PDAs, laptops, etc.) may coexist with very simple devices (sensors, RFID tags, etc.) and with an underlying network of sensors and storage elements. In this setting, im-

plementing efficient mechanisms for information retrieval is a challenging task. In fact, wireless channel impairments, differences in device capabilities, constraints related to the energy consumption, and node mobility are all aspects that may strongly influence the data dissemination phase.

While many current approaches and solutions for such ubiquitous networks are dedicated to specific applications, we aim for a generic and open approach for information exchange in such highly distributed environments. This means in particular that we do not assume knowledge about the actual destination of some piece of data in advance. For instance, some data read by a device passing by an RFID tag can be distributed proactively to other devices in the network. In this way, devices can obtain data without an end-to-end route to the source. We use distributed storage in the network, but also differ from centralized or infrastructure-based approaches, as we distribute the information opportunistically to storage devices in the network. While this approach cannot always guarantee optimal or full availability of all data (which may even be impossible in some scenarios), we argue that our approach is effective and does save the cost for a dedicated infrastructure.

In summary, our scenario is characterized by the following new technical challenges:

- The actual sources and sinks of information may not be known beforehand. The reason for this can be that the actual usage is not known beforehand but can also be due to mobility or sporadic communication.
- The network may not be fully connected at all times due to the mobility of the nodes as well as link or device failures. This will require to develop solutions that are robust to these impairments. Nevertheless, observe that node mobility may be exploited to deliver information to devices that, due to connectivity problems, are not directly attached to any communicating backbone.
- Our distributed system should work in the presence of

heterogeneous devices, with different capabilities in terms of transmission rates, processing, storage, and energy.

While several approaches aiming at distributing the data in a flexible manner in ad hoc networks exist, e.g., epidemic or probabilistic routing [1], [2], we argue that more efficient opportunistic strategies are needed to distribute information in this environment. We adopt here an approach where proactive data dissemination and reactive request mechanisms are combined to get the maximal benefits in terms of network performance and energy efficiency, while requiring a minimal degree of coordination among nodes.

For the dissemination mechanisms we consider a network coding strategy, which has proven to be very well suited to broadcast information to a large number of nodes [3]. With network coding, a node sends out coded packets consisting of (possibly linear) combinations of the information units available at that node. All received coded information is stored in a decoding matrix until sufficient information has been gathered to allow decoding. The key concept is that coding allows to distribute different combinations of the same information over many devices. As a consequence, communicating with a subset of sufficiently “information rich” devices allows to retrieve the desired information as long as the aggregate decoding matrix has full rank [4]. While there is considerable research on network coding for wireless networks [5] the performance of this technique in our scenario, with non-homogeneous data types and distributed asynchronous network operation, is not well understood.

The objective of the present paper is twofold. First, we give a detailed description of potential applications, by pointing out the related challenges and issues to be solved in order to realize our vision. Second, we introduce a novel approach for data distribution and retrieval and we show the validity of such an approach by means of simulation.

The rest of the paper is organized as follows: in Section II we present our reference application scenarios. In Section III, we specify the requirements. In Section IV we introduce a simple mathematical model to describe our network. In Section V we outline our data dissemination approach, whose performance evaluation is carried out in Section VI. In Section VII we discuss possible improvements to the proposed scheme which we leave for future research. In Section VIII we give a short review of the existing literature and in Section IX we conclude the paper by discussing directions for future research.

II. DESCRIPTION OF APPLICATION SCENARIOS

In the following, we present several sample application scenarios. Actual application scenarios help to point out re-

quirements, characteristics, and open issues of the problem that we want to address.

A. Local Information Exchange in Shopping Malls

As a first example we focus on a shopping mall where each user may want to retrieve information about the mall, news, current offers, and the products via distributed in-network communication.

Let us assume that each item of clothing in a shop is equipped with a tag (RFID). Tags can communicate information to passing devices, which are used as *data mules* in order to build our distributed network for information storage, processing and retrieval. With the network in place, users entering the mall could get general information about the shops. We can also think of users passing in front of a shop being automatically informed about the available offers. A crucial point in our communicating system consists of the cooperation and interaction among nearby devices. Note that the spatial correlation among information sources and consumers can and should be exploited while disseminating information. Also, users may query the system in order to retrieve specific information or send feedback to the information sources in the network about their interests. This may help in creating and maintaining a view on how data flows have to be distributed/re-directed. Some of these issues will be discussed in more depth later in this paper. Finally, privacy can be an issue here. On the other hand, if users distribute information on behalf of others anonymously, it is quite difficult to collect personal information about users.

B. Lost & Found: Opportunistic Object Management

Consider an office building where items like books or tools are widely used and distributed within the offices. In addition, several pieces of each item may be available and users may want to know the location of some of them. In classical infrastructure based solutions, a central database would be maintained which stores who is using each tool. Yet this database creates overhead and cannot answer simple queries like “Is there a Java book close to my office?”.

As an alternative approach, which motivates our research, we assume that items in the building are marked with intelligent tags which can be read through wireless communication (e.g. RFID tag readers). To avoid a dedicated infrastructure, we further assume that users have mobile devices equipped with tag readers. As users enter the communication range of these items, they can collect useful information about their location and subsequently spread it opportunistically to other devices. In this way, we exploit both mobility to transport data and locality of

the information to accumulate it. With such an approach for object management, users can send queries to neighboring devices to look for specific or matching items.

C. Information Exchange in an Entertainment Park

Consider an entertainment park where customers can collect information, e.g. take videos from the local attractions. This data can be distributed locally to inform others about the events and also to annotate and comment attractions. In addition, reviews or simple messages can be communicated to other users in the park by means of multi-hop communication. In this scenario, the sources and sinks of information are not known a priori, and data should be distributed in a suitable way such that users can access the content they desire. Note that data search may be difficult if little meta-data is available. Also, users may want to preview clips instead of getting full videos right away. Further, information may exhibit similarities and hence data compression/aggregation may be exploited to improve performance.

III. REQUIREMENTS AND RESEARCH AREAS

The network functionalities required by the above reference scenarios pose quite a few challenges due to the various features involved. Below we focus on the most important research issues.

A. Efficient Dissemination of Information

Since the communication system is highly dynamic, we need to design solutions which are robust to node mobility and link/device failures. This should be done carefully by minimizing data retrieval time from any point in the network and, at the same time, by taking advantage of mobility and heterogeneity of the devices. Different processing and storage capabilities translate into different network protocol constraints. Finally, due to the potentially large amount of information and the need to duplicate data for both increased robustness and reduced access time, the information should be carefully distributed.

B. Efficient Representation of the Information

While some devices may have considerable storage capacity (laptops or PDAs, for instance), storage may be a problem for limited devices (i.e., sensor tags). For this reason, we need to devise an intelligent and compact way to represent data which should be tunable in order to match the device capabilities.

C. Heuristics for Information Exchange

Each time two devices are within communication range of each other, we have a transmission opportunity or “contact”. The duration of a contact depends very much on device mobility. Our algorithms should decide, depending on the freshness of the information, its scope, the capabilities of the two communicating entities and the mobility of the devices, how much and which data should be exchanged between the two devices. Hence, the study of realistic mobility models is crucial to correctly test the proposed solutions.

D. Algorithms for Efficient Query Dissemination

Once the information has been stored in the network in the form of a distributed database, we design algorithms that allow one to query the system, such that data can be retrieved in a short time and with limited network resource consumption. This may be done, for example, by exploiting queries sent along progressively expanding rings. In this case we may think of optimizing the scope of each ring, the frequency of the search phases and so on. Some application scenarios could require stronger constraints concerning the timeliness of delivery or the reliability of the information. Also, efficient query dissemination can speed up the retrieval of the requested data.

E. Optimal Usage of Storage and Networking Resources

It is desirable to optimize the position of the various static network components, such as sensor nodes and storage elements. This can be seen as a backbone network which can be used to assist the distributed information handling. In many practical cases, sensors and storage elements cannot be optimally placed but they are randomly deployed within the network area. In these settings, it is interesting to investigate how to exploit the backbone in the most efficient way and how the performance worsens with respect to an optimal placement.

IV. NETWORK MODEL

In this section, we develop a simple model for our system. The scenario we try to model is kept as general as possible and can be tailored to reflect different types of applications.

Let us assume to have a set V of n nodes randomly placed in a rectangular area. Each of them is a potential source of information in that it can generate information itself (for instance in case of sensors) and can also store information generated by other nodes. Information data is classified into Q different categories: $\{1, 2, \dots, Q\}$. We assume that only a subset of the nodes are *information sources* (sources). Similarly, a subset of nodes act as

information consumers (or sinks) and are interested in getting the information provided by the sources.

Let S_i be the set containing the nodes that produce data of type i , with $|S_i| \leq n$. Furthermore, let D_i be the set containing the nodes interested in receiving data of type i . Each node $v \in D_i$ wishes to receive data from a subset S_{iv} of nodes in S_i . The following constraints have to be satisfied:

$$\bigcup_{i=1}^Q S_i \subseteq V, \bigcup_{i=1}^Q D_i \subseteq V, \bigcup_{v \in D_i} S_{iv} \subseteq S_i, 1 \leq i \leq Q \quad (1)$$

According to this, we can model a number of possible cases of interest. For example, taking $Q = 1$, we get the following cases:

- All to all: $S_1 = D_1 = V, \forall v \in D_1 S_{1v} = V$.
- Broadcast: $\exists s \in V : S_1 = s, D_1 = V, \forall v \in D_1 S_{1v} = s$.
- Multicast: $\exists s \in V : S_1 = s, D_1 \subset V, \forall v \in D_1 S_{1v} = s$.
- Unicast: $\exists s \in V : S_1 = s, D_1 = d, S_{1d} = s$.

We also include all intermediate cases. To further improve the described scenario we can introduce some spatial and temporal correlation among the positions of the sources and the consumers. Indeed, it may be reasonable to consider that most of the nodes interested in the reception of the same data type i are placed close to each other. Therefore, it is possible to enrich the model choosing the set D_i and S_{iv} according to their geographical location. Moreover, note that the subsets of information consumers and sources could be varied dynamically in time to reproduce the behavior of different users and applications.

V. TECHNICAL APPROACH

The focus of our work is on a fully distributed system which shall be built on top of a network with dynamically changing connectivity. We aim at designing a set of algorithms that will be at the same time autonomous and adaptable. As a central tool for the communication part, we take advantage of network coding to efficiently use the available network capacity.

As we observed before, one of the main challenges in our reference scenarios consists of gathering information without any a priori knowledge about its existence, its location or the features of the information consumers. There are two main approaches to do this. In a *proactive scheme*, sources actively distribute their data towards the sinks. An alternative approach is to use a *reactive scheme* where sinks broadcast their queries and the sources reply if they have the desired information.

We are interested in combining these two approaches and in investigating, under which conditions hybrid schemes

may outperform either of the two mechanisms. In general, proactive schemes perform better when the number of consumers is high, while reactive strategies may lead to better performance in the opposite case [6]. We observe due to the characteristics of network coding [4], it will provide more benefits in proactive solutions, where the data is disseminated via broadcast to all nodes.

In summary, we can distinguish two methods to disseminate the information:

- 1) *Proactive data dissemination* or, in short, *push schemes*: the information is disseminated by means of flooding. Here, network coding can be used for improved performance.
- 2) *Reactive data dissemination* or *pull schemes*: these mechanisms can be subdivided in two phases:
 - a) The *inquiry phase*: information consumers *query* the system searching for specific information.
 - b) The *reply phase*: nodes respond to the queries by sending the desired information to the consumers.

Note that the above mechanisms are complementary in nature but can be implemented in parallel. Indeed, our idea is to combine them over the time in order to obtain a unique hybrid scheme.

VI. EVALUATION OF DATA DISSEMINATION WITH NETWORK CODING

The aim of this study is to examine the number of data transmissions in the network when hybrid (proactive and reactive) data dissemination is implemented. Specifically, we are interested in evaluating the overhead introduced by several dissemination schemes by varying how much information is disseminated proactively vs. reactively in order to identify in which situations a hybrid approach, where both reactive and proactive phases are present, performs best.

A. Introduction to network coding

Network coding is an emerging technique to efficiently disseminate data. Unlike in classical store and forward mechanisms, information packets are encoded at intermediate nodes in the network and subsequently forwarded. Among other coding techniques, linear coding can be used and has been proven to be effective [7]. This is a very good finding as it allows to implement these techniques through simple operations and, for this reason, is very often used. In this paper we focus on linear coding due to its low complexity.

Consider the case where a node receives, as input, a given number of packets. These, according to the network

coding strategy, are linear combinations of some pieces of information that are to be disseminated through the network. The node receives the new data, checks whether it is possible to decode it and re-transmits a (possibly) new linear combination of the received/stored packets. This mechanism can be described by a theoretical model which is sketched in the following (see [7], [8] for further details). Let us focus on n nodes in the network and let $[x_1, \dots, x_n]^T$ be the source symbols (or in general packets) generated at these nodes. Symbols are defined over a finite field F_q of q elements. Each linear combination \mathbf{y} over F_q can be described as the product of a vector of coefficients (*coding vector*) and the vector of source symbols:

$$\mathbf{y} = \mathbf{a}^T \mathbf{x} = \begin{bmatrix} a_1 & \dots & a_n \end{bmatrix} \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix} \quad (2)$$

Note that in order to decode at any given node in a distributed and asynchronous manner, it is necessary to receive both the encoded information \mathbf{y} and the related coding vector \mathbf{a} . In [4] the authors propose to send the coding vector along with \mathbf{y} in the same packet. Each node v stores its source symbol and the information vectors it receives in a decoding matrix \mathbf{G}_v . A newly received packet is said to be innovative if its vector increases the rank of the matrix. Reception of non-innovative packets is simply ignored. A node v will transmit a linear combination that lies in the vector space of its decoding matrix \mathbf{G}_v . Once a node has received n linearly independent combinations (and hence \mathbf{G}_v has full rank) it is able to decode and retrieve the original symbols x_1, \dots, x_n .

In a multicast scenario, linear codes considerably improve the throughput of the system [7]. Further work, by Koetter and Médard [9], focuses on finding the coefficients of linear coding and decoding functions for a given network, by studying some of their properties as well as the time needed for their calculation. A further solution, proposed in [4] is to randomly choose (in F_q) the coding coefficients. The strength of random coding is that if the local coding vectors are generated randomly, and the symbols lie in a finite field of sufficient size, the decoding matrix \mathbf{G}_v has full rank with high probability. Most of the current work on distributed and asynchronous wireless ad hoc networks [10], [11] exploits the random coding technique as, while retaining most of the benefits, it is computationally efficient and does not require synchronization among nodes.

Finally, observe that when we do not have any information about the position of the data consumers, the best forwarding strategy would be to re-combine and re-broadcast all the received packets. This strategy, in fact, would ensure that all pieces of information will eventually

reach every node in the network, and this is actually the best that can be done in the absence of any *a priori* knowledge.

B. Hybrid Data Dissemination Scheme with Network Coding

As explained in Section V, we investigate an algorithm that consists of a push phase to proactively disseminate data and a subsequent pull phase to obtain missing information.

We introduce a so-called forwarding factor d , which controls how much information is disseminated in the push phase. The factor controls how much information is sent out upon reception of new information and thus controls, how far new information will propagate from the original source. If $d < 1$, for each received innovative packet, a new packet is generated from the matrix and sent out with probability d ; if $d \geq 1$, $\lfloor d \rfloor$ packets will be sent and a further packet is sent with probability $d - \lfloor d \rfloor$. Refer to [11] for a comprehensive treatment of this approach.

The lower the forwarding factor, the lower the overhead (in terms of number of transmissions) of the first phase, and the lower the fraction of packets that can be decoded at the individual nodes. To allow all nodes to decode, the pull-phase is used to request the missing information. The pull phase itself consists of two parts: first, nodes send out queries to request the missing data and second, replies are sent by nodes that have the missing data until all nodes are able to decode all the data.

A query packet is sent out by each node with missing information (i.e., a node whose decoding matrix \mathbf{G}_v does not have full rank). Queries are sent using expanding ring search. A time to live (TTL) field in the packet limits the scope of the search. In the first search round, a given sink sends the query to all one-hop neighbors (TTL= 1). On receiving the query, the neighbors check whether they can contribute innovative packets and if yes, they return them via unicast transmission. If, after the first round, the decoding matrix still has less than full rank, the consumer initiates a new search with the TTL field set to twice the previous value, until finally all data can be decoded.

The overall cost incurred by such an approach consists of the cost of the push phase, the cost to send the queries, and the cost to fetch the missing packets via unicast. In general, the cost to send the missing packets depends on the hop distance between the information producers and the information consumers.

We compare this approach with a flooding based scheme that uses the same phases. The only difference is, that flooding instead of network coding is used for the first phase and that replies to information queries are uncoded original packets instead of coded information. For the push phase of flooding, a packet that is received for the first

time is rebroadcasted with probability d , such that network coding and flooding incur roughly the same overhead for the same forwarding factor.

C. Simulation setup and results

We assume a static wireless network with $n = 100$ nodes randomly distributed within a $1250m \times 1250m$ area. Nodes have a radio range of 250m, resulting in approximately 12 neighbors per node. The MAC layer uses an idealized version of CSMA-CA with a random schedule. Time is slotted with slot length equal to the packet transmission time. All nodes are sources and each has one original packet to disseminate. To analyze the tradeoff between the two phases, we vary the number of receivers that are interested in these original packets (from all-to-all communication to all-to-one communication). We measure overall cost in terms of total number of transmissions for the different phases.

Figure 1(a) shows the combined cost for the two phases for both network coding and flooding as a function of the forwarding factor d . Here, all nodes are interested in receiving all the information and consequently using the push phase to disseminate everything is most efficient. For low forwarding factors, much of the information is missing and each node has to send many queries. The more information is disseminated, the lower the overall cost, since dissemination to everyone is less costly than individual searches for information. For very high forwarding factors, overall cost increases again, since almost all nodes receive all the information and it is less expensive to send out a one-hop query than to risk sending a lot of redundant information to nodes that could already decode. The minimum total overhead for both phases is achieved for a first phase forwarding factor $d = 0.2$ for network coding and for $d = 1.0$ for flooding. Due to the large number of different linear combinations that can be sent, a network coding scheme is much less likely to transmit non-innovative information to nodes that already have some of the information. This allows to achieve a delivery ratio close to 100% for a much lower overhead.

In contrast, if only one of the nodes is a receiver, proactive dissemination is much less desirable, as shown in Figure 1(b). The minimum total overhead is achieved for a very low forwarding factor, meaning that almost all information is provided in the second phase. As expected, queries for missing information sent out by only a single node are costly than distributing the information proactively to all nodes. Since the second phase is the same for the network coding approach as well as the flooding scheme, there is almost no difference in performance between the two.

While these two corner cases allow us to study the basic behavior of the algorithm, scenarios with an intermediate number of receivers are more realistic and more interesting to analyze.

We simulate the same scenario with the same number of nodes but with the number of receivers ranging from 1 to 100. The left graph of Figure 2 shows the forwarding factor that results in the minimum total overhead for different numbers of receivers, with the resulting total overhead shown in the right graph of Figure 2. As the number of sink nodes increases, the optimal forwarding factor increases and more information should be disseminated in the push phase. Interestingly, the cost savings of network coding during the push phase result in the same overhead for network coding when delivery data to all receivers as the flooding overhead to deliver data to around 5 receivers (where most of the overhead comes from the pull phase).

This leads us to several conclusions. First of all, network coding offers significant benefits over flooding in a network with multiple receivers. Disseminating data to *all* receivers with network coding is less costly than a combination of flooding and expanding ring search with unicast delivery *only* for sink nodes, even if only 5% of the nodes are actually interested in the data. Second, significant cost savings are possible when there are *very* few receivers by reducing the amount of information that is sent out proactively with network coding. This indicates, that in a heterogeneous environment as in the aforementioned application scenarios, the optimal combination of the two phases depends to a large degree on the receiver density and is specific to a network area. Clusters of receivers with similar interests should be served by a directed push of information that transports most of the required data, while individual isolated receivers should rather search for the information they require. To this end, a feedback phase is required that transports information about the type of interest in information in certain areas of the network back to forwarding nodes and the sources. We discuss design considerations for such a feedback phase in the next section.

VII. FEEDBACK-BASED PROACTIVE DATA DISSEMINATION

The role of this section is to discuss possible improvements to the previously presented dissemination schemes. In fact, information sources may have or acquire some knowledge about the consumers (their location, type of interests, or the number of nodes interested in receiving the data). Hence, they could concentrate their efforts in providing the information in the right direction. A total or partial knowledge about the location of the consumers could be achieved by considering a *feedback mechanism* as

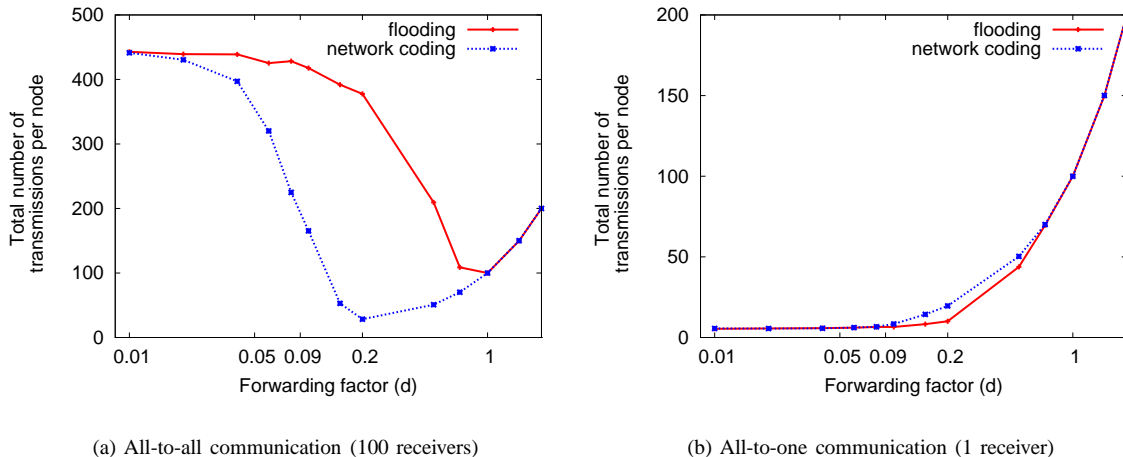


Fig. 1. Overall cost for push phase and pull phase to achieve 100% packet delivery ratio with network coding and flooding for different forwarding factors.




Fig. 2. Optimum forwarding factor (left) and corresponding overhead (right) with different numbers of sinks for network coding and flooding.

further step in the proactive scheme. In the following two subsections we investigate this aspect in more detail.

A. Enhancing the Proactive Scheme

In the basic scheme discussed above, the information is disseminated thanks to a proactive flooding phase, where network coding may be used for enhanced performance.

We however observe that in many practical situations not all devices will be interested in receiving all the information generated within the network. For instance, in our shopping mall example in Section II-A, users standing in front of a given shop will most likely be interested in retrieving information related to items of the same type of those sold in the closeby store. In particular, the probability that a given information will be requested from a device in a specific network location may depend on the location

itself. In fact, the interest on some specific items or, more generally, piece of information, may be influenced by the surroundings, as often happens in everyday life. For example, customers walking in proximity of a pizza place are often tempted to buy some food. This might happen even if they are not hungry and just because of the inviting smell. Hence, it might actually make sense to account for an *a priori* distribution of the probability that a given information will be requested from a given location in the network.

This concept can be further illustrated by looking at Fig. 3, where we refer again to our shopping mall example. A group of information sources stand in front of a shop selling electronics (referred in the figure to as *electronics store #1*), whereas more potential customers stand in front of two further stores which sell items of the same type.

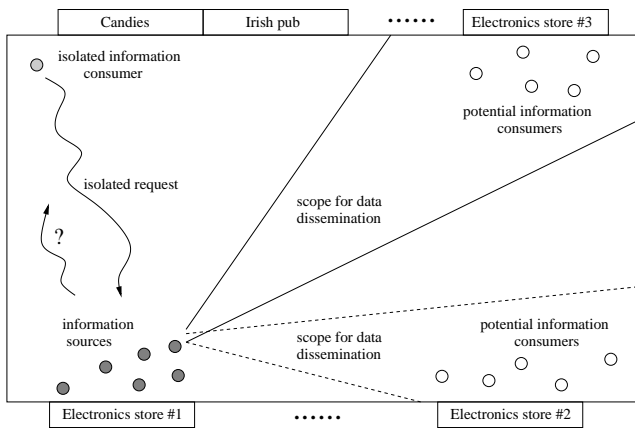


Fig. 3. Shopping mall example.

According to the above discussion, we may argue that the information generated by the nodes in close proximity of store #1 is likely relevant to the customers in close proximity of the remaining two stores selling electronics.

Motivated by the above arguments, we advocate the exploitation of some form of feedback to dynamically learn the portions of the network that will most probably contain users interested in a given type of information. We then plan to exploit such a knowledge within the proactive data dissemination by implementing a *controlled data dissemination* instead of pure flooding. In practice, we will try to balance the amount of information that will be sent in every portion of the network according to the received feedback.

The main advantage provided by such a form of controlled flooding consists of the following fact: in the absence of information regarding the position in the network of possible data consumers, the best we can do is to route the data randomly, that is, as soon as a node receives a packet it re-broadcasts it by re-combining (if network coding is used) with locally stored data. There is no reason not to broadcast the data as we actually do not know anything about the network. However, if we knew that the considered data are relevant for the nodes within a particular portion of the network only, then we could optimize the forwarding behavior at individual nodes. In particular, one might think of having rough indications of the portions of the network where the data is, *with high probability*, currently requested or, anyway, *is going to be requested*. If we had such information, we could think of properly re-broadcasting the data so as to achieve the following goals: 1) maximize, among the nodes interested in the data, or that are likely going to be interested in its reception, the number of devices that actually get the desired information 2) minimize the transmission overhead, which is equivalent

to minimizing the number of broadcast actions that will not serve any user. Henceforth, we may exploit the feedback to achieve the above goals by proactively avoiding to disseminate the data towards the portion of the network where such data is not needed. Of course, the approach is effective as long as we cut the right tradeoff between the savings enabled by controlled flooding and the extra overhead due to the additional *feedback* phase.

B. Feedback mechanism

Here, we briefly discuss the meaning of the feedback mechanism. This strategy consists of a form of feedback that we intend to propagate, from the possible points (devices) in the network where the information may be consumed, back to the points where the information is originated (information sources) or stored. The aim is to create a gradient (exploiting an appropriate system of coordinates) in order to drive the proactive dissemination phase towards the portions of the network where the information is most likely to be requested and/or is currently requested by a sufficiently large number of users. The feedback mechanism will provide us with the following information:

- *Triggering*: the feedback tells us that somewhere in the network somebody needs a given information (or may request it with high probability). Hence, the primary aim of the feedback is to provide information on the fact that somebody “wants the information” and that he/she “wants it now”. This will trigger a change in the forwarding and combining (network coding) rules.
- *Distance and location estimation*: the feedback can be exploited to retrieve information about the distance separating the nodes from the potential information consumers. This might be exploited to infer the costs that will be incurred in providing the wanted information.
- *Correlation*: we can exploit the correlation between subsequent feedback messages to infer, at a given node, which packets we should combine (and re-broadcast) if we intend to serve the consumers placed within a given area. With network coding in place as the communication mechanism, the idea is to broadcast the combination of packets which is most innovative for the region that we intend to serve.

A good method to create feedback messages may be to group the information composing our database in *a few* classes and propagate feedback messages for each class. Referring to Fig. 3, a possible class might be “electronics”, whereas other classes may be “clothes”, “food” and so on. In fact, in order not to overwhelm the network with feedback traffic, we should propagate as few messages as possible. We however think that even a rough view of

the network could considerably improve the performance in terms of diminished amount of *forward* information which is propagated without need. Besides, a controlled flooding technique might considerably help in decreasing the network interference and the resource consumption, as during the dissemination phase we do not transmit information where it is not needed. We stress that the feedback mechanism should be implemented in such a way to minimize the transmission overhead, which is intended as the energy required to transport the feedback itself. This may be achieved by aggressively exploiting data aggregation techniques along the reverse paths (consumers \rightarrow sources).

We finally note that as a reaction to an *interest* obtained with the feedback scheme, we might decide to change our forwarding strategy and explore new portions of the network that were previously unexplored. We also note that the feedback mechanism is optional. Should we not have it, then we would have to resort to the flooding. Our algorithm should be adaptable in this sense.

VIII. RELATED WORK

The “social network” concept has been attracting the attention of various researchers for several years and is now becoming a popular research topic. An interesting and recent work in this sense can be found in [12], where a wireless virtual social network is intended as a network that mimics the way in which people seek information. The work builds on the observation that people often are a great source of unique information which is likely location-specific, community-specific and/or time-specific. The authors in [12] define a specific infrastructure, called PeopleNet, aimed at propagating queries of a given type to users in specific geographic locations, which are referred to as the bazaars, via a fixed infrastructure. However, such an approach could be infeasible in highly dynamic networks, especially if the wanted information is in close proximity. In fact, in such a case searching for the information through the infrastructure would be inefficient.

Other studies address the problem of efficiently disseminating information or queries over a distributed network. Well known techniques to deal with message delivery without location information are *epidemic routing* [13] and *probabilistic routing* [2]. These strategies build on the same rationale: messages may have to be buffered for a certain time at intermediate devices, whose mobility is exploited to bring messages closer to their destinations by exchanging them between nodes as they come in touch. More in detail, in epidemic routing, nodes forward messages a certain number of times to a randomly chosen subset of neighbors. A more interesting study, which is presented in [1], addresses epidemic routing in partially-connected ad

hoc networks where a connected path from the source to the destination does not always exist. In probabilistic routing, instead, a node forwards the message to the neighbors with the higher probability to deliver data to the destination.

One of the main aspects to take into account in a distributed and heterogeneous system is the ability to acquire the desired data from any point in the network. To this end, one of the best known strategies consists of querying the system to disseminate requests for the wanted information. The standard approach to query processing is to flood queries to all nodes first and then construct a spanning tree where each node maintains a routing table of its parents [14]. Flooding can be pruned through adequate indexing structures in order to increase efficiency. An efficient indexing scheme is represented by the Geographic Hash Table, which maps identifiers and nodes in a properly defined metric space [15]. On this matter, we also cite significant work on multi-resolution distributed storage [16], [17], where wavelets are exploited to efficiently summarize data in a hierarchical structure.

Notably, Gao et al. present the concept of *fractionally cascaded information* to provide efficient algorithms and theoretical bounds for answering range queries in wireless sensor networks [18]. The principle they propose is that a device should know a fraction of the information from distant parts of the network, in an exponentially decaying fashion by distance.

In sharp contrast to the above methods, [19] assumes that it is possible to store all data in each node of the network so that when a user is interested in some information he can query any node and immediately retrieve it. However, this implies the usage of large buffer sizes and a massive pre-distribution of the information to all nodes. This becomes impractical as data traffic or size of the network increases.

In [20], the authors focus on the case where data collection queries can be posed anywhere in the network and require quick and reliable access to the data of interest, which is randomly scattered throughout the network. They investigate some connections between network coding and decentralized erasure codes.

We finally cite the recent work on optimal gossiping and routing in [21]. The authors of [21] present for the first time the gossip network model where travelers can obtain information about the state of dynamic networks by gossiping with peer travelers using ad hoc communication. Travelers can subsequently use the acquired information to refine their path and find the lowest cost route to their destination.

IX. CONCLUSIONS

Motivated by the current trends in wireless technologies and market opportunities, in this paper we presented

challenging scenarios consisting of distributed moving and uncoordinated wireless objects. Our primary interest is to design network protocols to efficiently build and maintain such distributed systems in an efficient manner. The challenges to be faced are manifold; just to cite a few, we need to cope with the wireless transmission channel, the subsequent impossibility of having fresh and accurate views of the network at every device and the limited energy reserve, which is typical of mobile devices.

Further, we proposed proactive, reactive and hybrid strategies to disseminate information in such environments. Network coding was the selected technique for data dissemination. Our first results for the basic algorithm confirm that network coding can introduce benefits to disseminate information in such a distributed system and show how it is possible to combine proactive and reactive approaches to improve performance. Also, the results reveal that network coding largely outperforms standard solutions which are based on the store and forward paradigm.

In our future research, we will integrate optimization algorithms (such as feedback mechanisms) which will be used to improve the network coding mechanism that we use in the proactive dissemination phase. Mobility and other requirements and open issues pointed out in this paper will be considered as well.

REFERENCES

- [1] A. Vahdat and D. Becker, "Epidemic routing for partially-connected ad hoc networks," in *Technical Report CS-2000-06*, July 2000.
- [2] A. Lindgren, A. Doria, and O. Schelen, "Probabilistic Routing in Intermittently Connected Networks," *ACM SIGMOBILE*, vol. 7, no. 3, 2003.
- [3] J. Widmer and J.-Y. L. Boudec, "Network Coding for Efficient Communication in Extreme Networks," in *ACM SIGCOMM 2005*, Philadelphia, PA, US, Aug. 2005.
- [4] P. A. Chou, T. Wu, and K. Jain, "Practical network coding," in *51st Allerton Conf. Communication, Control and Computing*, Oct. 2003.
- [5] R. Koetter, "Network coding bibliography," URL: <http://tesla.csl.uiuc.edu/~koetter/NWC/Bibliography.html>.
- [6] J. Heidemann, F. Silva, and D. Estrin, "Matching data dissemination algorithms to application requirements," in *ACM SenSys'03*, Nov. 2003.
- [7] S.-Y. R. Li, R. W. Yeung, and N. Cai, "Linear network Coding," *IEEE Trans. on Information Theory*, vol. 49, no. 2, Feb. 2003.
- [8] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung, "Network information flow," *IEEE Trans. on Information Theory*, vol. 46, no. 4, July 2000.
- [9] R. Koetter and M. Medard, "An algebraic approach to network coding," *IEEE/ACM Trans. on Networking*, vol. 49, no. 11, Nov. 2003.
- [10] J. Widmer, C. Fragouli, and J.-Y. L. Boudec, "Low-complexity energy-efficient broadcasting in wireless ad-hoc networks using network coding," in *Proc. IEEE Information Theory Workshop*, Oct. 2004.
- [11] C. Fragouli, J. Widmer, and J.-Y. L. Boudec, "A network coding approach to energy efficient broadcasting: from theory to practice," in *IEEE Infocom*, 2006.
- [12] M. Montani, V. Srinivasan, and P. S. Nuggehalli, "PeopleNet: Engineering a wireless virtual social network," in *Proc. of ACM MobiCom*, 2005.
- [13] P. T. Eugster, R. Guerraoui, A. M. Kermarrec, and L. Massoulié, "Epidemic information dissemination in distributed system," *Innovative Technology for Computing Professionals*, vol. 37, no. 5, May 2004.
- [14] Y. Yao and J. Gehrke, "Query processing for sensor networks," in *CIDR 2003*, 2003.
- [15] S. Ratnasamy, B. Karp, L. Yin, F. Yu, D. Estrin, R. Govindan, and S. Shenker, "Ght: a geographic hash table for data-centric storage," in *In workshop on Wireless sensor networks and applications*, 2002.
- [16] D. Ganesan, B. Greenstein, D. Perelyubskiy, D. Estrin, and J. Heidemann, "An evaluation of multiresolution storage for sensor networks," in *Proc. of the First ACM Conference on Embedded Networked Sensor Systems (SenSys03)*, 2003.
- [17] D. Ganesan, D. Estrin, and J. Heidemann, "Dimension: Why we need a new data handling architecture for sensor networks?" in *ACM Computer Communication Rev.*
- [18] J. Gao, L. J. Guibas, J. Hershberger, and L. Zhang, "Fractionally cascaded information in a sensor network," in *Symposium on Information Processing in Sensor Networks*, April 2004.
- [19] J. Kulik, W. R. Heinzelman, and H. Balakrishnan, "Negotiation-based protocols for disseminating information in wireless sensor networks," in *Wireless Networks*, vol. 8, no. 4, 2002, pp. 169-187.
- [20] A. G. Dimakis, V. Prabhakaran, and K. Ramchandran, "Ubiquitous access to distributed data in large-scale sensor networks through decentralized erasure codes," in *Symposium on Information Processing in Sensor Networks (IPSN '05)*, 2005.
- [21] Y. Shavit and A. Shay, "Optimal Routing in Gossip Networks," *IEEE Trans. Veh. Technol.*, vol. 54, no. 4, July 2005.