

Music Retrieval Over Wireless Ad-Hoc Networks

Ioannis Karydis, Alexandros Nanopoulos, Apostolos Papadopoulos, Dimitrios Katsaros, and Yannis Manolopoulos

Abstract—Wireless networks introduce new opportunities for music delivery. The trend of using mobile devices on wireless networks can significantly extend the recent change of paradigm in the model of music distribution by allowing mobile clients to search for audio music in a network of wireless mobile hosts. This paper introduces the application of content-based music information retrieval (CBMIR) in wireless ad-hoc networks. We investigate, for the first time in the literature, the challenges posed by the wireless medium and recognize the factors that require optimization. We propose novel techniques that attain a significant reduction in both response time and network traffic, compared to naive approaches. Extensive experimental results illustrate the appropriateness, effectiveness, and efficiency of the proposed method to this bandwidth-starving and volatility due to mobility and environment.

Index Terms—Content-based similarity, Mobile Ad-hoc Network (MANET), music information retrieval (MIR), peer to peer (P2P), reducing length representation, wireless ad-hoc network.

I. INTRODUCTION

A. Music Distribution Adopts a New Paradigm

IMAGINE listening to music through your enhanced pocket-sized ultralight device while jogging or resting in a park. A device that, apart from the ability to play prestored music like any MP3 player in a area that is not covered by wireless local area networks, can also search for and acquire music songs from other people's similar musical devices. This data exchange is attainable through the device's wireless connectivity equipment allowing for participation in ad-hoc networks, formed with similar devices being in close proximity. Although such a scenario may seem futuristic, it is not so distant.

Having already reached the end of an era for the traditional music distribution [25], the market model, as well as the buying

behavior of consumers, have been reformed by the development of technologies like MP3 (and the supporting applications for their distribution, e.g., *Apple's iTunes*, *iMusic* online music services) and the penetration of the World Wide Web. Peer-to-peer networks and the maturing distributed file sharing technology enable the dissemination of musical content in digital forms, permitting customers a ubiquitous reach to stored music files.

New opportunities for music delivery are additionally introduced by the widespread penetration of the wireless networks (wireless local area networks (LANs), general packet radio services (GPRSs), universal mobile telecommunications systems (UMTSs) [8]) such as the pioneering applications [26] supporting the distribution of MP3-based songs to 3G UMTS devices. These applications rely on the existence of a central server, which receives requests from and delivers audio files to the mobile clients. Though, aside from these single-hop infrastructured wireless networks, music delivery can also unfold over the emerging Mobile Ad-hoc Networks (MANETs). The wireless ad-hoc networks are peer-to-peer, multihop, mobile wireless networks, where information packets are transmitted in a store-and-forward fashion from source to destination, via intermediate nodes. Such networks are expected to give rise to scenarios like the one previously mentioned. The salient characteristics of these networks, i.e., dynamic topology, bandwidth-constrained communication links and energy-constraint operation, introduce significant design challenges.

In this paper, we focus on the following problem. We consider a number of mobile hosts that participate in a wireless ad-hoc network, where each host may store several audio musical pieces. Assume a user that wants to search in the wireless network, to find audio pieces that are similar to a given one. For instance, the user can provide an audio snippet (e.g., a musical piece excerpt) and query the network to find the peers that store similar pieces. As will be described in the following, the definition of similarity can be based on several features that have been developed (see Section IV-A) for content-based music information retrieval (CBMIR). It is important to note that the querying host does not have any prior knowledge of neither the qualifying music pieces nor the hosts' locations that contain them. This differentiates the current problem from existing ones that are interested just in identifying the hosts in a wireless ad-hoc network that contain a known musical piece. Moreover, the examined problem is complementary to the one of delivering streaming media (audio and video) [1] in wireless ad-hoc networks, since the latter does not involve any searching for similar musical pieces, and just focuses on transferring data from one host to another.

B. Requirements Set by the Wireless Medium

This research focuses on the development of methods for searching audio music by content in wireless ad-hoc networks,

Manuscript received January 20, 2006; revised October 18, 2006. First published June 27, 2008; last published July 16, 2008 (projected). This work was supported by the *HPAKAEITOS* and by a *FFET* grant in the context of the project "Data Management in Mobile Ad Hoc Networks" funded by the *ΠΥΘΑΓΟΡΑΣ II* national research program. The work of A. Nanopoulos was supported by an IKY postdoctoral scholarship. Preliminary results of this research appeared in the Proceedings of the International Symposium on Music Information Retrieval (ISMIR) in 2005. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. George Tzanetakis.

I. Karydis was with the Aristotle University of Thessaloniki, Thessaloniki 54124, Greece. He is now with the Department of Informatics, Ionian University, Corfu 49100, Greece (e-mail: karydis@delab.csd.auth.gr; karydis@csd.auth.gr).

A. Nanopoulos, A. Papadopoulos, and Y. Manolopoulos are with the Aristotle University of Thessaloniki, Thessaloniki 54124, Greece (e-mail: alex@delab.csd.auth.gr; ananopou@csd.auth.gr; apostol@delab.csd.auth.gr; apapadop@csd.auth.gr; manolopo@delab.csd.auth.gr; manolopo@csd.auth.gr).

D. Katsaros is with the University of Thessaly, Volos 38221, Greece (e-mail: dimitris@delab.csd.auth.gr; dkatsaro@csd.auth.gr).

Digital Object Identifier 10.1109/TASL.2008.925874

where the querier receives music excerpts matching to a posed query. As for the legal issues of transferring and reproducing the musical pieces found are concerned, analogous issues are being confronted in online music distribution over wired P2P networks, where ways to protect intellectual property are now maturing. CBMIR applications in wireless networks can, and must, adopt any such developments. Additionally, for the issue of reproduction, adequate techniques for preserving intellectual property do exist [12].¹

The searching procedure can benefit from the latest approaches for CBMIR in wired P2P networks (see Section II-A). Nevertheless, the combination of the characteristics of the wireless medium and of the audio-music data pose new and challenging requirements, which call for new solutions.

- 1) CBMIR methods for wired peer-to-peer (P2P) networks do not consider the continuous alteration of the network topology, which is inherent in wireless ad-hoc networks, since mobile hosts (MHs, the terms MH and peer are similar in this context and thus interchangeable) are moving and become in and out of range of the others continuously. One impact of this mobility is that selective propagation of the query among MHs, e.g., by using data indexing like distributed hash tables (DHTs), as proposed by [31], is not feasible. Additionally, the recall of the searching procedure is affected by the possibility of unsuccessful routing of the query, as well as the answers. Thus, new query propagation methods need be developed for wireless ad-hoc networks.
- 2) The need to reduce traffic, which results from the size of audio-music data. This is done by replacing the original query with a newly developed representation that utilizes novel, appropriate transcoding schemes. The requirement of traffic reduction is much more compelling in wireless ad-hoc networks, where the communication ability is usually assumed to be around 1 MB/s for relatively long distances.
- 3) In CBMIR over wired P2P networks, should a matching music excerpt be found, it can immediately be returned to the querying node, since the querier is directly accessible. In contrast, in wireless ad-hoc networks, the answers to the query have to be propagated back to the querier via the network. This requirement further burdens traffic, thus requiring optimization.

Existing methodologies in MANETs address the aforementioned issues in a limited extent. In particular, algorithms proposed for the problem of routing in MANETs consider neither the peculiarities of searching for CBMIR purposes nor the size of the transferred data, since music data are considerably larger than routing packets. These peculiarities include the need to search for similarity in a search space where similar content location is not known beforehand. As for the data, their size prohibits the dissemination for direct local retrieval. To our best knowledge, no existing approach has addressed all the aforementioned issues, collectively, in the field of MANETs.

¹The searching procedure, i.e., the subject of this work, does not reveal concerns about legal issues, since it only involves excerpts.

C. Contribution and Paper Organization

This paper introduces the application of CBMIR in wireless ad-hoc networks and investigates the challenges posed by the wireless medium in order to perform content-based MIR, as described in the paradigm of Section I-A. Accordingly, we propose novel techniques, which attain a significant reduction in both response time and network traffic, compared to naive approaches.

To address the requirements posed by the wireless medium, we propose the following techniques.

- 1) To fulfill the first requirement, we perform breadth-first searching over the wireless ad-hoc network using knowledge about neighboring MHs (obtained by probing neighborhood at specific time points). This approach can cope with mobility, maintain increased final recall, and constrain the drawbacks of flooding, e.g., excessive traffic due to multiple broadcasts (explained in Section II-B).
- 2) The second requirement is addressed by a technique that uses a concise, feature-based representation of the query with reducing length. The reducing-length representation (a.k.a transcoding) that we propose drastically degrades traffic, while reducing the computation performed at each MH as well.
- 3) The additional traffic produced by the third requirement is addressed by a twofold proposal. i) We propose policies to constraint the number of MHs involved for the propagation of the answers, by exploiting any MHs that were involved during the propagation of the query. (ii) We allow such MHs to prune the propagation of answers, based on a property of the previously described representation.

To our best knowledge, this work is the first to examine the issue of CBMIR in ad-hoc wireless networks. The contributions are 1) the introduction of the problem and the identification of the inherent requirements, 2) a novel algorithm that combines the aforesaid techniques and addresses the posed requirements, and 3) an extensive experimental evaluation, which illustrate the efficiency of the proposed methodology.

The rest of the paper is organized as follows. Section II describes background and related work. In Section III, we outline the proposed method, whereas Section IV describes features selection and the indexing method that we use. Section V provides a complete account of the searching algorithms and, subsequently, Section VI describes the proposed routing policies. Section VII presents and discusses the experimentation and results obtained. Finally, the paper is concluded in Section VIII.

II. BACKGROUND AND RELATED WORK

A. CBMIR in P2P Networks

Research related to the application of CBMIR in wired P2P networks is recent. In one of the first attempts, [32] presents four P2P models for CBMIR, which include centralized, decentralized, and hybrid categories. Another research based on a hybrid configuration is presented in [32], in which the authors propose a DHT-based system utilizing both manually specified attributes (artist, album, title, etc.) and extracted features in order to describe the musical content of a piece. The authors of [37] propose the utilization of the feature selection and extraction process that is described in [36] for CBMIR in a decentralized

unstructured P2P system. Moreover, although oriented towards a differentiated discipline, the work of [28] refers to audio retrieval in P2P networks. The principal target of this research is combating of unauthorized music file sharing in P2P networks. Finally, [15] investigated the problem of content-based searching for similar acoustic data over unstructured decentralized P2P networks, under the time-warping distance.²

In the present work, we deal with a wireless ad-hoc network, where two nodes can communicate only if in close proximity (in-range). As described, in this kind of network peers participate randomly and for short term, and when they do, they change frequently their location. These factors cause existing approaches, e.g., indexing, to become inapplicable.

B. Information Discovery/Provision in Wireless Mobile Ad-Hoc Networks

Ad-hoc networks are significantly different than wireless local area networks (WLANs), which are infrastructured, and wireless personal area networks (WPANs), e.g., Bluetooth, which are very short-range wireless networks (with a range around to 10 m. Although, different than WLANs and WPANs, MANETs are often implemented using WLANs or WPANs. Thus, the medium access control layer of the ad-hoc networks is commonly assumed to be than of WLANs or WPANs, providing, for instance, symbol rates at the range of 11 and up to 50 Mb/s. Though, these rates are achievable for ranges less than 70 m; for ranges between 110–130 m the rate is 1 Mb/s, whereas for distances longer than 100 m, the rates drop below 1 Mb/s. For this reason, almost all the studies involving transmissions at a range of 250 m or longer, assume a symbol rate between 500 Kb/s and 1 Mb/s.

In an ad-hoc network, when a source node desires to send a message to some destination node and does not already have a valid route to that node, it initiates a path discovery process to locate the destination. Nodes are identified by their IP address and maintain a broadcast ID, which is incremented after every route request they initiate. The broadcast ID together with the node's IP address, uniquely identify a route request. In the same manner, the transmitted data requests can be identified.

Remotely related to the topic of this paper is the issue of multicasting streaming media (audio/video) to MANETs (e.g., [9]) or unicasting audio to 3G UMTS devices [26]. These works though assume the existence of a central server/supplier, which provisions the mobile clients with multimedia data.

III. OUTLINE OF THE SEARCHING PROCEDURE

In this paper, music similarity is used in order to identify similar musical pieces to a query musical piece, in a network of mobile hosts (as described in Section I-A). The problem of finding similar music sequences in a MANET requires a searching procedure, which will detect MHs in the MANET that have similar sequences, find those sequences in the MHs, and return them

²Dynamic time warping (DTW) has been proposed as a more robust similarity measure to Euclidean distance, as it can express similarity between two time series even if they are out of phase in the time axis or they do not have the same length.




Fig. 1. Example of the (a) forward and (b) backward searching procedure.

back to the querier. The already described requirements of the wireless framework formulate the examined searching procedure in the following way.

- 1) There is no prior knowledge of the data MHs store, that is the querier has no knowledge of the location of the required data.
- 2) MHs that have qualifying sequences have to be reached in a way that addresses their mobility and minimizes traffic. Due to their relative positions and the preferred tolerance to traffic (see below), all such nodes may not be possible to reach.
- 3) At each reached MH, the qualifying sequences have to be detected by detaining the MHs, in terms of CPU cost, as little as possible.
- 4) Each qualifying sequence has to reach the querier in a way that reduces traffic. Notice that the answers may have to be routed back to the querier following paths different from those through which the MHs with qualifying sequences were reached, since intermediate MHs may have changed their position, and therefore be out of range. Due to this, every detected answer may not be possible to reach the querier.

An example is illustrated in Fig. 1. The querier is MH P_1 . During the forward phase [Fig. 1(a)], the query is received by MHs P_2 and P_3 . During the backward phase [Fig. 1(b)], answers can be directly returned by P_2 (still in range of P_1). Due to relative movement, P_3 is, now, out of range. Thus, its answers are routed through P_4 (previously out the range of P_1).

The searching procedure is initiated at the querying MH, aiming at detecting sequences in other MHs, which contain sequences whose similarity from the query sequence Q is within user-defined boundaries, a threshold ϵ . The definition of the distance measure is detailed in Section IV. Just for now, we can intuitively think of the distance as a measure of how dissimilar two music sequences are. The length of detected sequences is equal to the length of the query sequence Q .

To address traffic minimization, Q has to be transformed to a representation form, denoted as R , through which qualifying sequences are detected.

Due to this transformation, it is possible that false-positive results may appear. A false positive result is a result that appears to be a true result when comparing with the transformed representation, though, under the nontransformed query is not a real result. Moreover, R must present no false-negatives (real results that were missed due to the transformation). However, its particular implementation determines whether false-positives may be produced or if they will be completely avoided. Based on all the

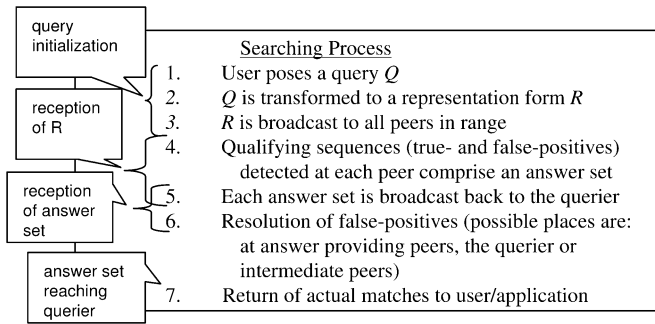


Fig. 2. Searching process and basic events.

aforementioned issues, an abstract scheme to describe the entire searching procedure consists of the steps depicted in Fig. 2, which are also summarized in four events.

To avoid duplicate effort, the procedure tags R with an ID (see Section II-B). This way, MHs that have already received it will perform no further action. Additionally, the propagation of R to the neighboring MHs is controlled by a parameter called h , which is a counter that is decreased at each receiving MH (denotes the available number of hops). Its initial value, at the querier, is equal to $MaxHop$. This value corresponds to the preferred tolerance to traffic and network reach/coverage. The propagation of answer sets (resulting from step 5) is handled similarly.

As already mentioned, the searching process consists of a forward and a backward phase. During the former, R is propagated and during the latter answers are routed back to the querier. The two phases are interleaved, since during the propagation of R by some MHs, other MHs are returning answers to the querier. The backward phase's volume mainly depends on the existence of answers and the number of false-positives, while the forward phase depends on the size of R , our coverage willingness as well as the network reachability. In general, the volume of information transferred during the backward phase is larger than that of the forward phase.

Having outlined the searching procedure, in the following sections we detail its parts. First, we elaborate on the features that can be selected for the formation of R . Next, we describe the acceleration of similarity searching within each MH by using indexing. Based on these, we next describe two searching algorithms, which follow different choices with respect to the formation of R . Finally, we present methods to improve the backward phase.

IV. FEATURES AND INDEXING

A. Features for CBMIR

One of the main challenges in MIR is the choice of representation of the musical information within the system. As the sequences of acoustic music objects tend to be quite large in size, they are commonly described by a set of features. Numerous standpoints exist on what features to retain and on how to select these features [3], [22]–[24] [33], [34]. Meaningful features help in the effective representation of the objects and enable the use of indexing schemes for efficient query processing. The most typically encountered features for the acoustic representation

are produced by time analysis, spectral analysis, and wavelet analysis.

In this paper, we do not concentrate on devising new features. Instead, we are interested in a methodology for the searching procedure. Our methodology is able to embrace any high-performance feature extraction procedure. Accordingly, we apply a feature extraction process based on the wavelet transform. Wavelet transforms provide a simple but yet efficient representation of audio by taking into consideration both nonuniform frequency resolution and impulsive characteristics, as shown by [4], [18], [19].

The wavelet transform has long been used in image and signal processing while its use in information retrieval and data mining has been extensive [21]. A complete survey on wavelet application in data mining can be found in [18]. In general terms, the wavelet transform is a tool that provides quality time and frequency resolution, while dividing up data, functions, or operators into different frequency components and then studying each component with a resolution matched to its scale.

Wavelets present numerous favorable properties in contrast to other type of analyses. Among them, lie the efficient computation complexity, the vanishing moments that support denoising and dimensionality reduction while focusing on the most important information, the compact support that guarantees the localization of the wavelet, the decorrelated coefficients that enable the reduction of complex processes of time domain into simpler in the wavelet domain and the support to the Parseval's theorem. In addition, wavelets present a multiresolution property that leads to hierarchical representations and manipulations of the objects treated.

The previously mentioned merits of the wavelet transform corroborate the use of wavelets on music. The low computation complexity assists the already burdened process by the large size of the musical data. The vanishing moments and their denoising capability cope with the noise introduced in musical recordings by the ambient sounds, during recording. The compact support allows locally altered musical pieces to retain their overall similarity, while the multiresolution adheres to the perception model of the ear, according to which the perception of both large-scale quantities and small-scale events, rely upon the multiresolution capability of the ear [4].

More particularly, we consider the Haar wavelet transformation for its simple incremental computation, its capability concerning the capture of time dependant properties of data and overall multiresolution representation of signals [7] as well as for the incorporation of the previously mentioned properties. However, our approach can easily be extended to other types of wavelet transforms.

B. Indexing Within Peers

To facilitate the searching within peers, we use the following approach. In a peer, each original audio sequence is transformed to a number of multidimensional points. We use a sliding window of length n over the sequence and apply discrete wavelet transform (DWT) to the contents of each window, producing n coefficients per window. Therefore, each audio sequence produces a set of n -dimensional points in the feature space. Since n depends on the query length and, thus, takes

relatively large values (e.g., 64 K), in order to efficiently index them in the feature space, we select only the first d dimensions from each point (in our experiments we used $d = 64$). This procedure dramatically reduces both the size of the index and the number of dimensions without affecting much the quality of the index. The reason for the latter is the merit of DWT to concentrate the energy of the sequence in the first few coefficients. However, false-positives are possible and thus require resolution.

Most importantly, it has been proven by [6] that no false dismissals are introduced when using only the d first coefficients (due to Parseval's theorem). Notice that this property is proven in [6] for the Euclidean distance. Although this distance measure is simple, it is known to have several advantages, as it has been illustrated by [17]. Nevertheless, the proposed methodology does not decisively depend on the particular features and distance measure, which are used herein following simplicity as well as computation efficiency reasons.

To speed-up the retrieval, for each sequence the collection of the resulting d -dimensional points is organized in minimum bounding rectangles (MBRs), which are, then, stored in an R^* -tree [2]. Answering to query, the root is initially retrieved and its entries that intersect the query are only further examined recursively until reaching a leaf. All nonintersecting nodes are not included in the search. Therefore, when searching for similar subsequences, we first retrieve candidates from the R^* -tree. We rank the candidates so as to process the most promising ones first (we observed that this saves a lot of CPU time) and then, those candidates are examined against the provided query representation. When the latter is reduced (as in the case of transcoding that will be explained), false-positives are still possible. Nevertheless, their number is significantly reduced. More details about indexing can be found in [14].

V. SEARCHING ALGORITHMS

In this section, we describe the two algorithms that implement the searching procedure. The first is based on simple choices concerning the representation R of the query sequence and its propagation during the forward and backward phases. The second (proposed) is based on more advanced choices with respect to the latter issues.

A. Algorithm Based on Maximal Query Representation

A simplistic approach for the representation R is to set it identical to the query sequence. The advantage is that no false-positives occur, since when a possible match has been found by index probing, it can be immediately tested against the query itself (i.e., R). Thus, no false-positives will be included in the answer-sets, which could negatively impact the backward-phase traffic, as they would be propagated to the querier just to find that they are not actual matches. We have to note that, to be able to perform index probing (i.e., to avoid sequential searching at each MH), a small number of DWT coefficients are included in R as well. However, their size is negligible compared to the size of the query sequence.

The resulting algorithm is denoted as ML (full *Maximum* representation with *Local* resolution at MHs). ML is summarized

- **Query initialization** The querier assigns to R the entire query sequence (plus the few query coefficients) and propagates (broadcasts) it to all its neighbors.
- **Reception of R** Upon the reception of R , each MH P probes its indexes, resolves the false-positives, and produces a list of results (only true-positives). The answer-set is propagated back to the querier, by broadcasting it to *all* the neighbors of P (backward phase). Accordingly, should there be available h , R is conveyed to all P 's neighboring MHs (forward phase).
- **Reception of an answer-set** Each MH P , that is not the querier, receiving an answer-set, continues the propagation (backward phase) to *all* its neighboring MHs as long as there is available h .
- **An answer-set reaches the querier** When an answer-set reaches the querier, then the results are immediately presented to the user.

Fig. 3. ML algorithm.

in Fig. 3 according to the actions performed for each occurring event (see Fig. 2).

Although ML manages to control the traffic during the backward phase (due to the elimination of false-positives), this comes at the cost of excessive traffic during the forward phase. This is due to the representation R that is propagated during the forward phase, which is equal to the entire query. For large query sequences this causes prohibitive forward traffic. Evidently, there emerges a tradeoff between the two contrasting phases. What is, therefore, needed is a method that will balance the traffic between the two phases, aiming at overall improvement.

Another issue on which ML makes a simplistic choice is the selection of the neighboring MHs to which the answer-set is propagated during the backward phase. When handling the second and third events, ML selects all neighbors for this purpose, thus resorting to plain flooding. This simplistic selection can significantly impact the backward traffic. To overcome the problem, we need to devise policies for the selective routing of the answer-sets. That is, we want to select only those nodes that are more promising to satisfy the receipt of the answers, thus significantly reducing backward traffic without reducing the chances of the answer-sets to reach the querier.

B. Algorithm Based on Reduced Query Representation and Transcoding

In Section V-A it was made clear that there is a tradeoff between the forward and backward traffic. ML focuses only on the improvement of backward traffic and incurs high forward traffic. In this section, we present a new algorithm, which has a twofold objective. The first is to produce a representation R that achieves a balance between the two phases and minimizes the overall traffic. The second is to develop selective routing policies for the propagation of the answer-sets, leading to significant reduction of the backward traffic.

The first objective is confronted by setting R between the two extremes cases: 1) the minimum possible representation with only the d DWT coefficients that are required for the local index-searching (minimizing forward traffic), and 2) the maximum possible representation with all n elements in the

query sequence itself (eliminating the burden of false-positives in terms of computation and backward traffic). Therefore, between the two extremes, R can consist of the l greater DWT coefficients, where $d \leq l \leq n$. Notice that this type of representation generalizes the two extreme cases: by setting $l = d$, R becomes identical to the first 1) case; in contrast, by setting $l = n$, R becomes identical to the second 2) case, because the n DWT coefficients are equivalent to the n elements of the query sequence (due to the Parseval's theorem).³ As described in Section IV, a number l of the greater DWT coefficients can effectively capture the energy of the music sequence and reduce the number of false-positives. The result is that, compared to the second 2) case, the forward traffic is expected to be smaller, because $l \leq n$. Compared to the first 1) case, the backward traffic is expected to be smaller too, due to the number of false-positives being significantly reduced, since $d \leq l$.

The tuning of l , however, is difficult, because it depends on several factors, like the topology of the MANET, which are changeable. For this reason, we follow a different approach. Initially, l is assigned a large value (see Section VII for its tuning) and this value is monotonically reduced during the propagation of R in the forward phase. This technique can be thought of as a *transcoding* scheme, as it involves sequences with varying number of DWT coefficients that correspond to varying approximations of the initial query sequence. The transcoding scheme:

- Keeps forward traffic low, as the size of R is reducing at each stage of the forward phase propagation.
- Reduces backward traffic by letting the MHs involved in the forward phase to cache the transcoded representation and, during the backward phase, to use it for early resolving false-positives, before they reach the querier. The problem of caching depends on several network parameters. This problem is independent to our approach, while effective solutions can be found in [10] and [16]. In our experiments, we found that by simply caching the representations for a small, fixed amount of time, adequate performance is attained.
- Reduces the processing (CPU) time at each MH, as the cost of resolving false-positives at each MH depends on the size of R .

The reduction is performed by getting l values according to an inverse sigmoid function [Fig. 4(b)]. Due to the shape of this function, the immediate neighborhood of the querier, which can provide results faster, receives a larger R , whereas the burden posed on MHs that are far is appreciably smaller. Also, this way we control the exponential growth of traffic that results by plain broadcasting of full-size representation. An example is depicted in Fig. 4(a). P_1 is the querier and P_4 is the node that starts propagating the answer-set. The MHs in the path from P_1 to P_4 are depicted gray shaded, and they are annotated with the size of R that reaches them (P_1 starts with 10 K DWT coefficients). Fig. 4(b) illustrates that these sizes are reducing, following an inverse sigmoid function. During the backward phase, starting




Fig. 4. Example of the searching procedure.

from P_4 , MHs P_3 and P_5 can be reached (depicted with dashed arrows). The cached representation in P_3 can help to resolve possible false-positives in the answer-set. The reason is that, in P_4 , the false-positives were examined against a smaller R than the one in P_3 . In contrast, P_5 was not in the path and thus cannot resolve any false-positives.

Henceforth, the size of the initial query representation is given as a factor (denoted as I) of the complete query size, whereas the slope of the inverse sigmoid function is controlled by a parameter denoted as α (higher values of α produce a steeper slope).

Regarding the second objective, we do not follow the simplistic approach of ML, which propagates the answer-sets to all neighbors. In contrast, during the forward phase, as it is typical in any dynamic source routing protocol [11], each MH that receives R , additionally receives the ID of all MHs that were used in the path from the querier to it. These IDs can be maintained along with R with minimal cost (only some bytes). When an MH starts propagating answer-sets, it selects among its current neighbors those that will propagate the answer-set (not all of them). To make this selection, it applies a policy that focuses on the neighbors that were included in the path from the querier to it. Since several such policies can be developed, in the next section we elaborate further on them. All the policies, despite their differences, emphasize on selecting neighboring MHs that were included in the path. The reason is that the cached representations that these nodes maintain can resolve false-positives during the backward phase. Therefore, traffic is substantially reduced. More details will be given in Section VI.

The algorithm that combines all the aforementioned characteristics is denoted as RT (querying by *Reduced representation with Transcoding*) and is illustrated in Fig. 5. The handling of success or failure is treated similarly to standard routing MANET protocols employing a time-to-leave (TTL)-like policy[35].

VI. ROUTING POLICIES FOR THE BACKWARD PHASE

In this section, we describe three policies for routing the answer-sets in the backward phase. The first two policies (global and local counter, described in [5]) are based on existing methods, whereas the third one (critical mass) is novel. As mentioned, all policies try to select nodes that were included in the path during the forward phase. Nevertheless, the backward phase cannot be based only on such nodes. Due to the mobility of MHs, it may be impossible to reach the querier unless other

³In the case of ML, we could have R to consist of all the n DWT coefficients. However, we choose the n sequence elements in the time domain just to avoid the computation of the inverse DWT, since in our case the time domain presents a smaller storage requirement as the data values are in range 0–255.

- **Query initialization** The querier sets R equal to a sample with an initial size (parameter) plus the query coefficients, and propagates (broadcasts) it to all its neighbors.
- **Reception of R** Upon the reception of R , each MH P probes its indexes, resolves as many false-positives as possible based on the received query sample of R , and produces a list of results. The answer-set is propagated back to the querier, by following the described policy for the backward phase. Accordingly, should there be available h , R 's size is reduced, and the reduced R is conveyed to all P 's neighboring MHs (forward phase).
- **Reception of an answer-set** When a MH receives a reply, it checks if it can resolve any false-positives. This is true should it have received (if any) a representation that was larger than the one that the sequences in the answer set were examined previously (i.e., at the sending MH). After any possible pruning, as long as there is available h , the answer-set is routed backwards following a policy.
- **An answer-set reaches the querier** When an answer-set reaches the querier, initially any remaining false-positives require resolution, and then the results are presented to the user.

Fig. 5. RT algorithm.




Fig. 6. Example of propagation in a MANET. (a) Forward phase. (b) Backward phase.

MHs (not included in the path) are additionally involved. The objective of all policies is to control the number of involved MHs so as to reduce backward traffic. These policies constitute a hybrid approach between probabilistic broadcasting, where the broadcasting decision is completely local to each mobile host and the deterministic broadcasting which relies on the discovery of some form of connected dominating set [20].

A. Global and Local Counter Policies

To clarify the description of the first two policies, consider the example of Fig. 6(a), which depicts the path from MH P_1 to MH P_6 , which was followed in the forward phase. Fig. 6(b) depicts the routing of the answer-set from P_6 back to P_1 . Comparing the two phases, several MHs have changed their location, others have switched off, and some new ones have become reachable. The MHs that are depicted grey color are the ones that were included in the forward path too, whereas the rest are new ones that were involved only in the backward phase.

With the *global-counter* (GC) policy, when an MH (P_6 in the example) starts propagating an answer-set, it tags the answer-set with a maximum number of retransmissions, h , equal to the length of the forward path plus an extra value e . In the example, the length (number of edges) is equal to 5. Let $e = 1$ and $h = 6$. GC tries to find among the neighbors the one that was its predecessor in the path. In the example, at P_6 GC tries to find P_5 . If this MH is reachable, then it is the only one selected to propagate the answer-set, and h is decreased by one. The same procedure is applied recursively. At P_5 , GC tries to find P_4 . If




Fig. 7. Example of relative locations of MHs in (a) forward and (b) backward phase.

P_4 is not reachable, as it is now the case, then GC propagates the answer-set to all neighboring MHs (broadcasting to P_7 and P_8), and each of them receives an h value decreased by one. Next, unless an MH in the forward path has been reached, GC continues by broadcasting to all neighbors. At each propagation of the answer-set, h is decreased by one, thus actually acting as a decreasing global counter. If an MH from the forward path is reached at any point again, then, as previously, GC tries to find its predecessor. In the example, P_3 is a node from the path, which has been reached with h equal to 2. Its predecessor is P_2 , which then propagates (as h is 1) the answer-set to P_1 and the procedure terminates.

In summary, when selecting the MHs to route back the answer-set, GC tries to follow the MHs included in the forward path. However, to overcome problems from the alteration of the MANET (like the disappearance of P_4 in this example), it allows an amount of discrepancy by resorting to broadcasting. To control the discrepancy, and thus the backward traffic, it uses the value of e . Notice that with a very large e , GC resorts to broadcasting for a very large number of times, thus becoming equivalent to the simplistic policy used by the ML searching algorithm. In contrast, with a very small e , the querier may not become reachable, especially when the MANET changes very fast.

A variation of GC works as follows. After a discrepancy, when an MH from the path has been reached again, we reset h to its initial value. In the previous example, when P_3 is reached again, available hop is reset to 6 (initial value). Thus, h acts as a decreasing local counter, because it is reset independently at several MHs. For this reason, this policy is denoted as *local-counter* (LC). Its objective is to increase the probability of reaching the querier, by rewarding the identification of the forward path. Nevertheless, this can increase the backward traffic.

B. Critical Mass Policy

With the *critical-mass* (CM) policy, if at least a number, denoted as *critical-mass factor* (CMF), of the current neighbors was in the forward path, we select them as the only ones to propagate the answer-set. If their number is less than CMF, then we additionally select randomly some of the current neighbors (not in the path) in order to have at least CMF MHs to propagate the answer-set. In contrast, if their number is larger than CMF, then they are all selected. For example, consider the case in Fig. 7. Fig. 7(a) depicts the forward phase, whereas Fig. 7(b) presents the backward case. As shown, during the backward phase some

MHs have now relocated. Let CMF be 2. When P_4 starts propagating the answer-set, it first selects P_3 , because it belongs to the forward path. Since this is the only such MH and CMF is 2, it also selects P_5 at random among the other reachable MHs.

The nodes that were selected at random in order to fulfill CMF, are still provided with the path of the MH that initiated the propagation of the answer-set (for the previous example, P_5 that is selected by P_4 , will also know the path from P_1 to P_4). This way, due to mobility, it is possible for such nodes during the backward phase to find neighbors that appear in the forwarded path (in the same example, P_5 finds P_2 that was in the path). Therefore, the impact of such randomly selected MHs on the proposed policy may be kept at a moderate level.

The CM policy differs from GC and LC in the following aspects. 1) It does not search predecessors in the path, as it focuses on identifying MHs that were in the forward path, regardless of their order (i.e., not searching for the previous node). This makes CM more flexible to the changes in the MANET. 2) It never resorts to broadcasting to all neighbors. At worst case, the number of randomly selected MHs is equal to CMF. This attains better control of the backward traffic. Due to the aforementioned characteristics, CM is expected to outperform GC and LC, as will be shown in Section VII.

VII. PERFORMANCE EVALUATION

A. Simulation Configuration

In this section, we provide an experimental comparison of the three described content-based audio retrieval algorithms. The performance of the algorithms was compared through simulation. The settings of the simulation were as follows. The mobile ad-hoc network had 100 nodes. We used 300 real acoustic sequences, which correspond to various pop songs. The average duration was about 5 min. To account for the fact that songs (especially the popular ones) are common in several nodes, we replicated each sequence to a number of MHs (default value equals to four). The aforementioned settings correspond to a realistic scenario for a MANET, like the one described in the Introduction. Accordingly, the average number of sequences per node was 12, a quantity that is quite reasonable for the state-of-the-art MP3 cell phones [27] and PDAs, both of which support the latest memory cards.

Regarding the simulation of mobility, we based our experiments on the Generate_Spatio_Temporal_Data (GSTD) simulator, as presented in [29], which considers hosts moving freely in a 2-D area. We used a squared area with side equal to 4000 m, whereas the transmission/reception range of each MH was set to a 500-m radius. Different degrees of velocity were selected for the moving MHs, adjusted by parameters of the GSTD, but due to lack of space we present results only for the average walking speed of a human (5 Km/h). Additionally, to account for the fact that mobile devices may enter doze mode (power-safe status where the device is out of network), we take each time for the MHs a doze-mode probability, with default value equal to 0.1 (that is, at each time unit a MH is out of network with probability 0.1).

Regarding CM, the default value for CMF was 10% of the number of neighbors at each MH, whereas the default initial

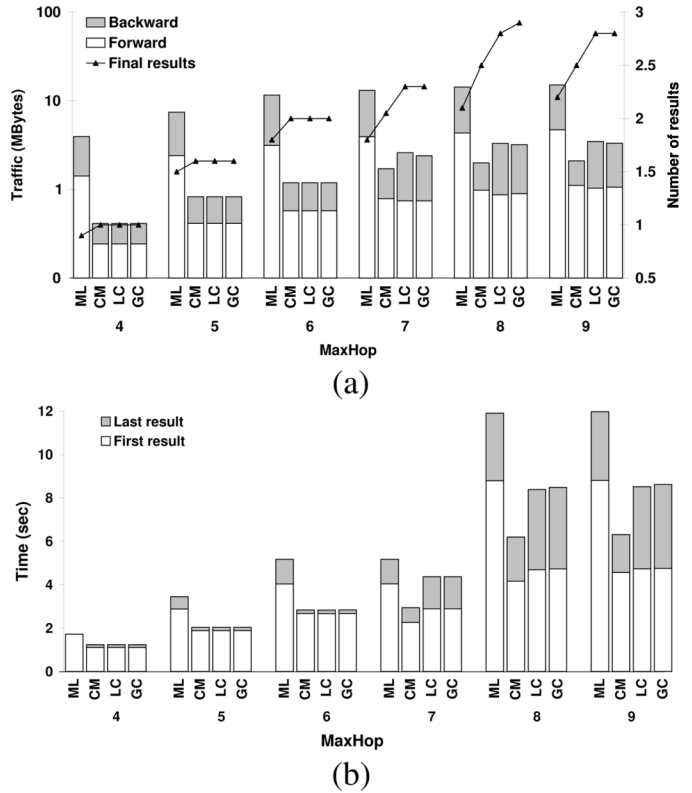


Fig. 8. Traffic, number, and time of results versus MaxHop.

sample size was 10% of the query sequence's size. For GC and LC policies, the additional value e added to MaxHop is set to 2 (we tried other similar values with no significant improvement). For all algorithms, the default value of ϵ was 0, and the default MaxHop was set to 5. Henceforth, when parameter values are not specified, we assume the default values.

The evaluation metrics are the average traffic (measured in megabytes) that each query incurs, the number of results obtained and the time the first and last result were discovered (the time of the first result is a useful measure, since users may terminate searching early). The results on time reflect the perceived latency required for the response to the querier. In contrast, total traffic reflects the load posed to the network in order to provide responses. Thus, the two factors require separate consideration.

B. Experimental Results

In our first experiment, we examined the traffic against MaxHop. The results are illustrated in Fig. 8(a). The forward and backward traffics are depicted separately, whereas their addition (height of bars) gives the total traffic. As expected, ML produces the highest forward traffic in all cases (due to maximal query representation), whereas the forward traffic of CM, LC, and GC are about the same. Regarding backward traffic, as described, ML attains a decreased number of returning results. However, due to the absence of an efficient backward routing policy, this advantage is invalidated. The rest approaches, considerably improve backward traffic, with CM performing better for MaxHop greater than seven. From this result it becomes obvious that, although the backward phase is in general more demanding for all algorithms, due to the reduction of

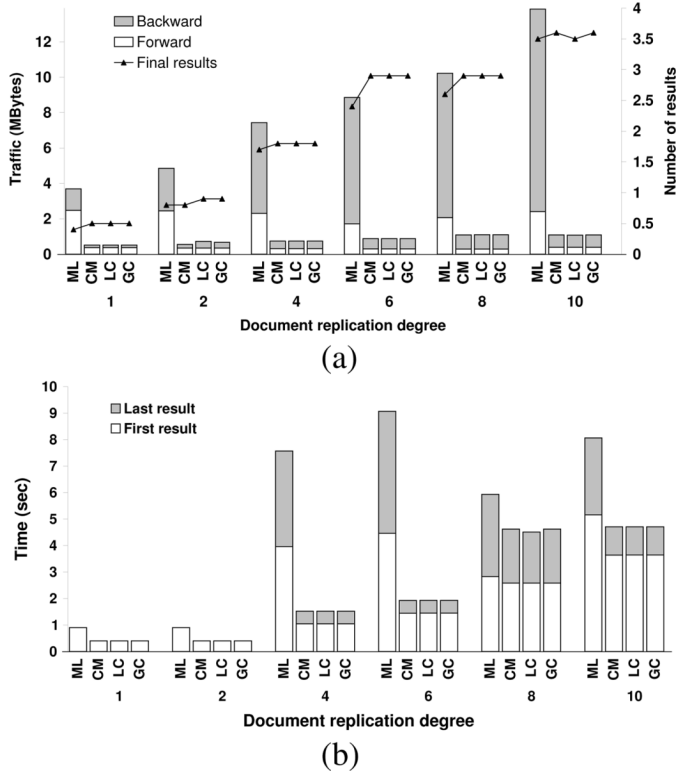


Fig. 9. Traffic, number, and time of results versus document replication degree.

backward traffic attained by CM, LC, and GC, the requirement of optimizing the forward phase, is fair. Additionally notice that the number of results (depicted in Fig. 8(a) with a solid line) obtained by ML are less than the results obtained by CM, LC, and GC and although the difference is small, there is a clear trend.

This result can be further clarified by the results on time of the first and last results, which are depicted in Fig. 8(b). In this figure, the height of the bars correspond to the time of last result, whereas the time of first result is depicted separately as its fraction. As expected, increase in available MaxHop produces longer times, since more MHs are examined. In all cases, the increase in time is far more steep for ML, while CM presents an advantage over LC and GC.

Next, we examined the impact of the document replication degree on the traffic and the time of the first and last result. The former experiment [given in Fig. 9(a)] shows that an increase in the replication degree has a clear impact on both the number of results and the backward traffic. This is especially true for ML. Once again, the number of results obtained by ML are slightly less than for the other. Regarding times of first and last result [Fig. 9(b)], an analogous behavior is observed.

The following experiment considered the traffic produced by the MH doze-mode probability (Fig. 10). It is quite clear that for increased values of probability, the network becomes less connected, thus leading in decrease of the results returning to the querier. ML is clearly outperformed, whereas the others perform about the same. What is more, the increase in doze-mode probability leads to a decrease in traffic, since the diminished

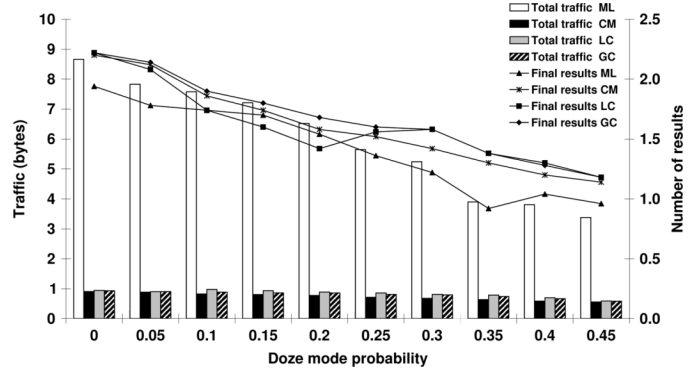


Fig. 10. Traffic and number of results versus doze mode probability.

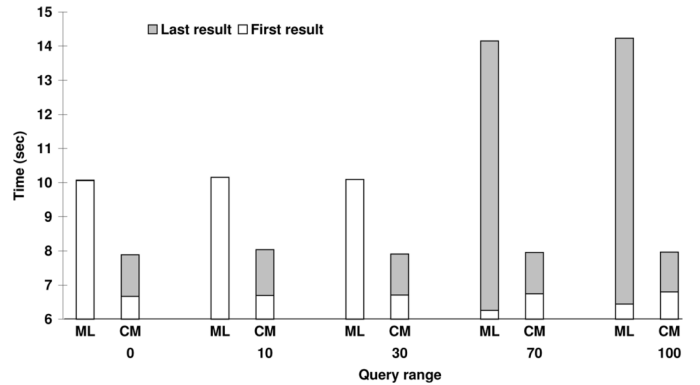


Fig. 11. Traffic versus query range (ϵ).

connectivity of the graph prohibits both the discovery of results and the propagation of any found.

Next, we examine the impact of query range ϵ . Fig. 11 shows the results for traffic with respect to ϵ . Since CM, LC, and GC perform similarly, to improve clarity we only include the results for the former. As ϵ increases, more results are found and, thus, backward traffic increases too (forward traffic is unaffected). However, the increase is much more obvious for ML, whereas CM, due the effectiveness of the policy for the backward phase, has a very smooth increase.

We also tested the sensitivity of the CM algorithm against I and α parameters (as formally described in Section V-B). The results are illustrated in Fig. 12(a). For both I and α parameters, the performance of ML remains totally unaffected and is only included for comparison purposes. As far as the I parameter is concerned, backward traffic is unaffected. As expected, forward traffic increases with increasing sample size. For the examined range of values, the reduction in traffic is not combined with a change in the number of found matches, which are similar in the order of decimal values, and thus omitted. In contrast, the examined values for the α parameter resulted in small differences between the approaches.

Finally, we examined the sensitivity of CM against CMF (the others are not affected by CMF). The traffic and number of results of CM, for varying CMF values, are depicted in Fig. 12(b). When CMF is high, the effectiveness of the policy for the backward phase is limited, since most MHs are selected at random by this policy. Thus, the resulting backward traffic is high (forward




Fig. 12. (a) Traffic versus I and α parameters. (b) Traffic and number of results versus CMF

traffic is not affected). Notice that, for the examined range of CMF values, the reduction in traffic slightly affects the number of found matches (the difference between the results for the extreme values of CMF are only in the order of decimal values). On the other hand, the increase in the results comes at the cost of higher traffic. Conclusively, relatively small CMF values are sufficient.

VIII. CONCLUSION

In this paper, we introduce the application of CBMIR application in wireless ad-hoc networks. We recognize the new challenges posed by this type of networks. To address them, we propose a novel algorithm, which is based on a twofold optimization: 1) the use of query representation with reducing length, and 2) a selective policies for the routing of answers, which performs additional pruning of traffic. The combination of these factors attains significant reduction in both response times and traffic. This is verified through extensive experimental results, which illustrate the suitability of the proposed method.

Concluding, we have to mention that the examined context does not depend on the specific features and distance measure, since it can be used in combination of several other ones, as long as they allow for a reducing-length representation.

In future work, we will examine other features. Additionally, we intend to extend our approach to accommodate musical genre querying as well. The key idea is that, based on a annotated querying feature set (such as the features described in [30]), the querier can identify similar genre audio data within a mobile ad-hoc network.

REFERENCES

- [1] L. Baochun and K. H. Wang, "NonStop: Continuous multimedia streaming in wireless ad-hoc networks with node mobility," *IEEE J. Sel. Areas Commun.*, vol. 21, no. 10, pp. 1627–1641, Oct. 2003.
- [2] N. Beckmann, H. P. Kriegel, and B. Seeger, "The R*-tree: An efficient and robust method for points and rectangles," in *Proc. ACM Int. Conf. Manage. Data (SIGMOD)*, 1990, pp. 322–331.
- [3] S. Bray and G. Tzanetakis, "Distributed audio feature extraction for music," in *Proc. Int. Symp. Music Inf. Retrieval (ISMIR)*, 2005, pp. 434–437.
- [4] A. P. C. Roads, S. Pope, and G. D. Poli, Eds., *Musical Signal Processing*. Lisse, The Netherlands: Royal Swets & Zeitlinger, 1997.
- [5] R. Castaneda, S. R. Das, and M. K. Marina, "Query localization techniques for on-demand routing protocols in ad-hoc networks," *ACM/Kluwer Wireless Netw.*, vol. 8, no. 2/3, pp. 137–151, 2002.
- [6] K. Chan and A.-C. Fu, "Efficient time series matching by wavelets," in *Proc. IEEE Int. Conf. Data Eng. (ICDE)*, 1999, pp. 126–133.
- [7] K.-P. Chan, A.-C. Fu, and C. Yu, "Haar wavelets for efficient similarity search of time-series: With and without time warping," *IEEE Trans. Knowledge Data Eng.*, vol. 15, no. 3, pp. 686–705, May/Jun. 2003.
- [8] J. DeVriendt, P. Laine, C. Lerouge, and X. Xu, "Mobile network evolution: A revolution on the move," *IEEE Commun. Mag.*, vol. 40, no. 4, pp. 104–111, Apr. 2002.
- [9] A. Dutta, J. Chennikara, W. Chen, O. Altintas, and H. Schulzrinne, "Multicasting streaming media to mobile users," *IEEE Commun. Mag.*, vol. 41, no. 10, pp. 81–89, Oct. 2003.
- [10] Y. Fang, Z. J. Haas, B. Liang, and Y.-B. Lin, "TTL prediction schemes and the effects of inter-update time distribution on wireless data access," *ACM/Kluwer Wireless Netw.*, vol. 10, no. 5, pp. 607–619, 2004.
- [11] D. Johnson and D. Maltz, T. Imielinski and H. Korth, Eds., "Dynamic source routing in ad-hoc wireless networks," in *Mobile Computing*. Norwell, MA: Kluwer, 1996, vol. 353, pp. 153–181.
- [12] T. Kalker, D. H. J. Epema, P. H. Hartel, R. L. Legendijk, and M. van Steen, "Music2share—Copyright-compliant music sharing in P2P systems," *Proc. IEEE*, vol. 92, no. 6, pp. 961–970, Jun. 2004.
- [13] I. Karydis, A. Nanopoulos, A. N. Papadopoulos, D. Katsaros, and Y. Manolopoulos, "Content-based music information retrieval in wireless ad-hoc networks," in *Proc. Int. Symp. Music Inf. Retrieval (ISMIR)*, 2005, pp. 137–144.
- [14] I. Karydis, A. Nanopoulos, A. N. Papadopoulos, and Y. Manolopoulos, "Audio indexing for efficient music information retrieval," in *Proc. IEEE Int. Conf. Multi Media Modeling (MUM)*, 2005, pp. 22–29.
- [15] I. Karydis, "Music retrieval in P2P networks under the warping distance," in *Proc. Int. Conf. Enterprise Inf. Syst. (ICEIS)*, 2005, pp. 100–107.
- [16] D. Katsaros and Y. Manolopoulos, "Web caching in broadcast mobile wireless environments," *IEEE Internet Comput. Mag.*, vol. 8, no. 3, pp. 37–45, May/Jun. 2004.
- [17] E. Keogh and S. Kasetty, "On the need for time series data mining benchmarks: A survey and empirical demonstration," in *Proc. ACM Int. Conf. Knowledge Discovery and Data Mining (SIGKDD)*, 2002, pp. 102–111.
- [18] T. Li, Q. Li, S. Zhu, and M. Ogihara, "A survey on wavelet applications in data mining," *SIGKDD Explorations*, vol. 4, no. 2, pp. 49–68, 2002.
- [19] T. Li, M. Ogihara, and Q. Li, "A comparative study on content-based music genre classification," in *Proc. ACM Int. Conf. Res. Develop. Inf. Retrieval (SIGIR)*, 2003, pp. 282–289.
- [20] W. Lou and J. Wu, I. Maghoub and M. Ilyas, Eds., "Broadcasting in ad-hoc networks using neighbor designating," in *Handbook of Mobile Computing*. Boca Raton, FL: CRC, 2004.
- [21] M. K. Mandal, T. Aboulnasr, and S. Panchanathan, "Fast wavelet histogram techniques for image indexing," *Comput. Vision Image Understanding*, vol. 75, no. 1–2, pp. 99–110, 1999.
- [22] D. McEnnis, C. McKay, I. Fujinaga, and P. Depalle, "Jaudio: A feature extraction library," in *Proc. Int. Symp. Music Inf. Retrieval (ISMIR)*, 2005, pp. 600–603.
- [23] M. Muller, F. Kurth, and M. Clausen, "Audio matching via chroma-based statistical features," in *Proc. Int. Symp. Music Inf. Retrieval (ISMIR)*, 2005, pp. 288–295.
- [24] F. Pachet and A. Zils, "Automatic extraction of music descriptors from acoustic signals," in *Proc. Int. Symp. Music Inf. Retrieval (ISMIR)*, 2004.
- [25] G. P. Premkumar, "Alternative distribution strategies for digital music," *Commun. ACM*, vol. 46, no. 9, pp. 89–95, 2003.

- [26] M. Rocchetti, P. Salomoni, V. Ghini, and S. Ferretti, "Bringing the wireless Internet to UMTS devices: A case study with music distribution," *Multimedia Tools and Applicat.*, vol. 25, no. 2, pp. 217–251, 2005.
- [27] "Mobile phones with memory cards," SanDisk, Aug. 2005 [Online]. Available: http://www.sandisk.com/pdf/retail/Mobile_Phones_with_Memory_Cards.pdf
- [28] P. Shrestha and T. Kalker, "Audio fingerprinting in peer-to-peer networks," in *Proc. Int. Symp. Music Inf. Retrieval (ISMIR)*, 2004, pp. 341–344.
- [29] Y. Theodoridis, J. Silva, and M. Nascimento, "On the generation of spatiotemporal datasets," in *Proc. Symp. Large Spatial Databases (SSD)*, 1999, pp. 147–164.
- [30] G. Tzanetakis and P. Cook, "Musical genre classification of audio signals," *IEEE Trans. Speech Audio Process.*, vol. 10, no. 5, pp. 293–302, Jul. 2002.
- [31] G. Tzanetakis, J. Gao, and P. Steenkiste, "A scalable peer-to-peer system for music information retrieval," *Comput. Music J.*, vol. 28, no. 2, pp. 24–33, 2004.
- [32] C. Wang, J. Li, and S. Shi, "A kind of content-based music information retrieval method in a peer-to-peer environment," in *Proc. Int. Symp. Music Inf. Retrieval (ISMIR)*, 2002, pp. 178–186.
- [33] A. Wiczorkowska, "Musical sound classification based on wavelet analysis," *Fundamenta Informaticae*, vol. 47, no. 1/2, pp. 175–188, 2001.
- [34] A. Wiczorkowska and Z. Ras, "Audio content description in sound databases," *Web Intelligence: Res. Develop.*, pp. 175–183, 2001.
- [35] B. L. Y. Fang, Z. Haas, and Y. Lin, "TTL prediction schemes and effects of inter-update time distribution on wireless data access," *ACM/Kluwer Wireless Netw.*, vol. 10, no. 5, pp. 607–619, 2004.
- [36] C. Yang, "Efficient acoustic index for music retrieval with various degrees of similarity," in *Proc. ACM Int. Conf. MultiMedia (MM)*, 2002, pp. 584–591.
- [37] C. Yang, "Peer-to-peer architecture for content-based music retrieval on acoustic data," in *Proc. ACM Int. Conf. World Wide Web (WWW)*, 2003, pp. 376–383.



Ioannis Karydis received the B.Eng. degree in engineering science and technology from Brunel University, London, U.K., in 2000, the M.Sc. degree in advanced methods in computer science from Queen Mary University, London, in 2001, and the Ph.D. degree in mining and retrieval methods for acoustic and symbolic music data from the Aristotle University of Thessaloniki, Thessaloniki, Greece, in 2006.

From June 2006 to February 2007, he was a Postdoctoral Researcher in music databases at the Aristotle University of Thessaloniki. He is now an

Adjunct Lecturer at the Department of Informatics, Ionian University, Corfu, Greece. His research interests include music databases, music information retrieval (indexing and searching), music mining, music object representation, and cultural informatics.



Alexandros Nanopoulos was born in Craiova, Romania, in 1974. He received the M.S. degree from the Department of Informatics, Aristotle University of Thessaloniki, Thessaloniki, Greece, in 1996, and the Ph.D. degree from the same institute in 2003.

Currently, he is a Lecturer in the Department of Informatics, Aristotle University of Thessaloniki. He is coauthor of more than 60 articles in international journals and conferences. He has also coauthor the monograph *Advanced Signature Techniques for Multimedia and Web Applications* and *R-trees: Theory*


and Applications. His research interests include data mining, web information retrieval, and spatial database indexing.



Apostolos N. Papadopoulos was born in Eleftheroupolis, Greece, in 1971. He received his five-year Diploma degree in computer engineering and informatics from the University of Patras, Patras, Greece, in 1994 and the Ph.D. degree from Aristotle University of Thessaloniki, Thessaloniki, Greece, in 2000.

Currently, he is a Lecturer in the Department of Informatics of Aristotle University of Thessaloniki. He has published several research papers in journals and proceedings of international conferences. From March 1998 to August 1998, he was a Visitor Researcher at INRIA Research Center, Rocquencourt, France, performing research in spatial databases. His research interests include spatial and spatio-temporal databases, data stream processing, data mining, and information retrieval. His research work has over 270 citations in scientific journals and conference proceedings. He has served as a track cochair of ACM SAC DTTA (Database Technologies Techniques and Applications) Track 2005, 2006, 2007, and 2008.

Dr. Papadopoulos is a member of the Technical Chamber of Greece.



Dimitrios Katsaros received the B.Sc. and Ph.D. degrees in informatics from the Aristotle University of Thessaloniki, Thessaloniki, Greece, in 1997 and 2004, respectively. The subject of his dissertation was "Information Dispersal in the Wireline and Wireless Web."

Currently, he is a Lecturer in the University of Thessaly, Volos, Greece. He is coeditor of the book *Wireless Information Highways* (IRM, 2005). His research interests include Web and Internet (particularly caching, replication, prefetching, and content

delivery), mobile and pervasive computing (especially data management and delivery), and data mining.



Yannis Manolopoulos was born in Thessaloniki, Greece, in 1957. He received the B.Eng. degree in electrical engineering in 1981 and the Ph.D. degree in computer engineering in 1986, both from the Aristotle University of Thessaloniki, Thessaloniki.

Currently, he is a Professor at the Department of Informatics, Aristotle University of Thessaloniki. He has been with the Department of Computer Science, University of Toronto, Toronto, ON, Canada, the Department of Computer Science, University of Maryland, College Park, and the Department of

Computer Science, University of Cyprus. He has published about 200 papers in refereed scientific journals and conference proceedings. He is coauthor of the following books: *Advanced Database Indexing* (Kluwer, 1999) and *Advanced Signature Indexing for Multimedia and Web Applications* (Kluwer, 2003), as well as *Nearest Neighbor Search: a Database Perspective* (Springer, 2004) and *R-trees: Theory and Applications* (Springer, 2005). His published work has received over 1700 citations from over 450 institutional groups. His research interests include databases, data mining, Web and geographical information systems, bibliometrics/webometrics, and performance evaluation of storage subsystems.

Dr. Manolopoulos served/serves as General/PC Chair/Cochair of the Eighth National Computer Conference (2001), the Sixth ADBIS Conference (2002), the Fifth WDAS Workshop (2003), the Eighth SSTD Symposium (2003), the First Balkan Conference in Informatics (2003), the Sixteenth SSDBM Conference (2004), the Eighth ICEIS Conference (2006), and the Tenth ADBIS Conference (2006).