A Preliminary Design for a Privacy-friendly Free P2P Media File Distribution System

Ron G van Schyndel

School of Computer Science and Information Technology
RMIT University PO Box 2476v, Melbourne 3000, Victoria
ron.vanschyndel@ieee.org,
WWW home page: http://www.cs.rmit.edu.au/~ronvs/

Abstract. In most P2P business models, in which users purchase the media, it is necessary to securely identify the user in order to facilitate payment. In this paper, we present a technique for allowing the widespread sharing of certain media formats including music using a method that keeps track of media possession and other marketing information, but in a way that does not require user identification. For the user, the main attraction of this scheme is that their identity is not a requirement, usage of reduced-quality media within this system is free and that extended media search is facilitated as an attraction to remain within the system. The content creators and distributors are compensated by this system by them having access to potentially large-scale actual usage and music trading statistics. The preliminary system design presented here, can cleanly coexist with a full-quality music purchase business model, also described briefly.

1 Introduction

In recent years, peer-to-peer networking has become an efficient and particularly scalable mechanism for widely distributing large media and data files. One application class that has proliferated on P2P networks is file-sharing and in particular, media-file sharing. This has been largely due to the popularisation of programs such as the original Napster [16], KaZaa [12] and Altnet [2], Gnutella [8], eDonkey [6] and many others. However, until 2003, most of this file sharing activity has been in breach of copyright law.

In 2003 Apple introduced iTunes [4] - a DRM controlled legal music-sharing system that allowed music-sharing under certain strict conditions. Since then, many other systems have emerged (Walmart [18], Musicmatch [15] - with Microsoft being the most recent to date [14] - each with minor variations to the same set of conditions).

All the above legal schemes have a subscription or pay-per-song business model. This model forces the media objects to assume the status of a commodity - a user must buy the object. This necessarily requires the user to provide their identity so as to secure payment.

In this paper, we describe, in an overview form, an initial design for a possible mechanism that entirely bypasses the need for user identity in securing access to music files available within this scheme.

We present here an approach where the economic value is not so much embedded in the shared media objects themselves, but in their relationship with each other - their context. While it is relatively easy to 'hijack' a media object such as a song which has intrinsic value and then pirate its value, it may not be so easy to hijack context.

This view directly benefits music publishers, which can use the context information for marketing. This marketing information is highly sought-after, especially if it is demonstrably representative of large population segments.

The information is not collected directly from a user's PC, but from the ISP through which the content is delivered.

1.1 A User-centred Approach

The design of this system commences from the user's view of the media in a business model which is a little more involved than a simple fee-for-service or fee-for-product.

In this system, a user collects media tags and uses them to facilitate media playing and exchange. A typical media-active user may end up collecting a large number of media objects and corresponding tags. The meta-information stored in the tags - to which the user can add extra free-form information, mostly for personal use - will allow the user to manipulate the media collection in ways not commonly available before. The user is thus encouraged to use such tags – their key value being not "Who has a particular tag?", but "Which sets of tags are commonly found together, and how can this information about the mix of music each user possesses be leveraged towards directed marketing?".

When the user registers their music, some of the tag information is embedded as meta data into the media file itself, however if a user has a media file with meta data and a such a tag, then where appropriate, the tag data locally overrides the embedded meta data. The degree to which tags information is embedded into the media is under user control - indeed, the value of the system to the user is that the user has full control over exposure of usage and other data.

Release of this information is encouraged, and rewarded through the improved searchability compared to current peer-to-peer systems, and through the use of 'freebies' and other marketing schemes. However, since no identifying information *ever* leaves the user, he or she is always free to 'turn off the tap' at any time with no penalty other than the loss of those extra benefits.

The value to the content provider/marketer/publisher is the possession statistics and the potential size of the population from which these are drawn.

Importantly, this system also allows for the discovery and trial of new music and other media content without risk to the user. It primarily uses the traditional word-of-mouth recommendation from friends, but now users can trial the whole song at low-quality. This is in contrast to systems such as iTunes, where content must first be purchased, or as with the 30-second samples used by Amazon

[3], (invariably the samples are unrepresentative of the song (it seems), as they were automatically produced). It is also possible to implement push-style media distribution without detracting from the rest of the system.

In contrast to some honour-based schemes such as Gnutella [8] and KaZaa [12], the 'free-riding' phenomenon [1] is avoided since all users are equal under this system in terms of the information transferred.

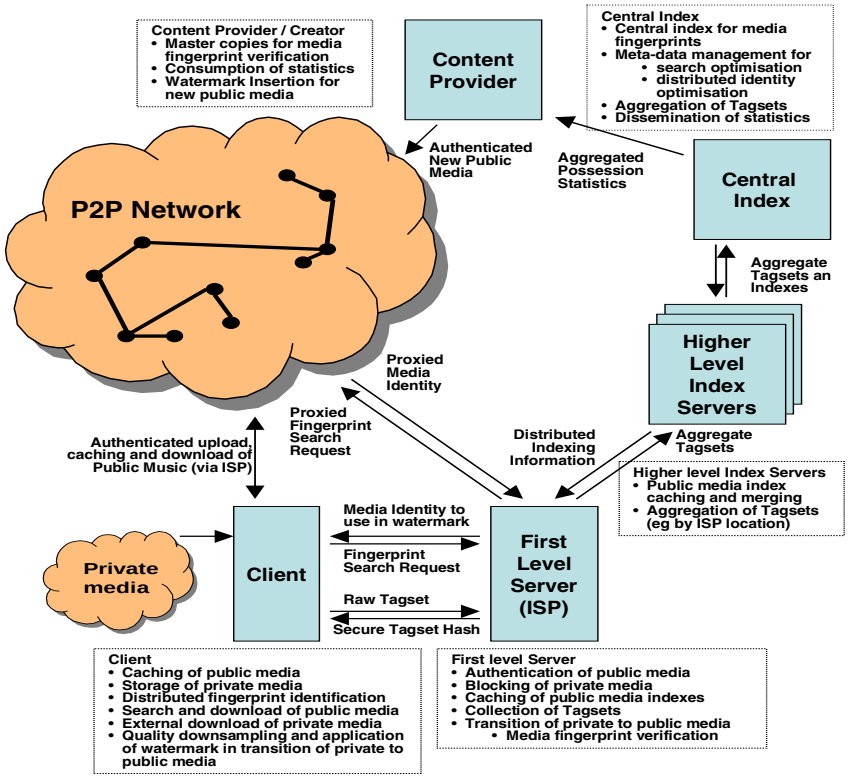

Fig. 1. Overview of the System

This system also fits with a full-purchase business model called Music2Share described by Kalker et al [13]. A user would use our system for free low and medium-quality media exchange in exchange for statistical information about which compatible media a user possesses and how that changes over time. This can be done without user identity being required. Those users who want guaranteed high or maximum quality media may use the purchase model suggested by Music2Share. Again, our system and Music2Share are mutually compatible.

2 Preliminary Design

We describe a system that is composed of a hierarchical set of media indexing servers, leading up to one or a small set of central servers which maintain a central media registry or index.

Figure 1 shows an overview of the system showing major protocol paths.

2.1 Media Types

There are three different types of media involved in this system. These correspond roughly to the similarly-named media types described by Kalker [13] in Music2Share (M2S).

Registered Public media: This is media that was obtained by the client from our system, or has been authenticated and registered by the client. It is unencrypted but bandlimited to a quality less than VHS or CD. A user who wants better quality must purchase the media (e.g. from M2S). Public media contains a watermark (such as in [19]), which is used for easy identification and contains the data required to render good quality audio/video. The high capacity watermark contains the lower and higher frequency components removed by the bandwidth limiting process of registration. A compatible player will decode the watermark and render in good quality once the song is registered.

Unregistered public media: This is public media which has not been registered by the client. Such media will play in a very bandlimited manner on compatible media players until registered, or on ordinary players which are not part of the system. Such media, while rendering badly, will still produce recognizable renditions. For example, music will render as if played through a telephone - it will still be recognisable and identifiable.

Encrypted public media: This is full bandwidth media of verifiable quality as per the M2S system. It is encrypted and the compatible media player can be instructed to silently pass the media to a M2S player when encountered. Note that this media type is not formally part of our system, but is compatible with it. It cannot be transferred as part of our system since user identity is required at all times.

Private media: This is media supplied by the user which has not been authenticated in any way. This media may come from anywhere, and the client is free to use it in any way. However private media will not be transmitted via our system and the client will not be able to take advantage of any indexing.

Authentication involves verifying the identity of the media from its 'fingerprint' (as with [13], then creating a down-sampled or bandlimited version of the media and inserting the watermark to create a public version of the media. This public version is then stored locally for P2P use.

Registration involves obtaining the identity and feature data of any new public media and combining it with the data collected from other public media in the client's possession (called a tagset), and then transferring that information to the media server. The server returns a secure hash of the tagset, and registration is complete.

Table 1. Media Types

Media Type	Quality	Usage	Bandwidth	Video size
Unregistered Public	Poor	Preview	6kHz	80x50
Registered Public	Public	Good Normal	15kHz	320x240
Private, Encrypted Public	Excellent	Premium	> 44kHz	full-screen

2.2 Media Index servers and Possession Statistics

The Media index servers - typically operated by ISP's - are specially authenticated servers that form a federated information resource on tag movements, and this information, in the form of aggregated statistics files can be used by Copyright agencies to determine possession-based royalty payments.

The media publishers pay for the privilege of obtaining the statistical data for sets of songs from ISP's running an index server. The ISP can also derive income from advertising and various marketing-driven games like an online radio or TV station. Sets of tags (as playlists) can be delivered as part of a promotion.

It is expected that over time there would be a natural merging of server operations like the above over a few large ISP's, but ISP franchise arrangements could be used to distribute this load.

2.3 A Hierarchical Peer-to-peer Approach

Many existing DRM-based music distribution systems such as iTunes [4] are based on a central server. This has the obvious disadvantage of high network load, poor scalability and net bandwidth bottlenecks, although this is not an insurmountable problem as shown by the architecture of Google [9]. One major difference is the average transaction size, which for Google is relatively small.

In terms of music sharing, the principal advantage of a centralised server model as in the original Napster [16] network model is the indexing. A central index can allow rapid match and complex search criteria as opposed to dynamically maintaining distributed indexes over many peers

By contrast, a P2P network has the advantages of being naturally scalable - distributing the network load and bandwidth, and providing a fault-tolerant manner of operation - any node may fail without major impact to the rest, as opposed to the single point of failure implicit in any client-server architecture. But a fully P2P system also has some distinct disadvantages: quality of service is uneven; slow propagation of information about new nodes; and sub-optimal searchability, especially with fuzzy search terms.

We therefore propose a compromise between a central server model and a fully decentralised model: a hierarchical P2P (HP2P) network using distributed hash tables along the lines of Chord and TOPLUS [7].

A HP2P model combines the scalability advantages of P2P with the centralised indexing of a traditional client-server model. Indeed, a number of distributed indexing models exist, each with different search and network complex-

ities (SearchTools [17], Collab [5]). Providing hierarchies allows for considerable search optimisation and may be sufficient to make fuzzy searches feasible.

Except for the largest ISP's, the first-level hierarchy would most likely encapsulate all the users for a particular ISP. In the second hierarchical level, multiple ISP-based song indexes would be combined. In other words, while the file-sharing would remain strictly P2P, the indexing would be more hierarchical.

If an ISP is distributed geographically, then there may be virtue in splitting on that basis. In this case, dynamic indexing may allow certain local meanings or interpretations to match more strongly, if the search request was made locally.

2.4 Sharing Music

Music is shared between users by exchanging the tags or the media (usually the former). The system will then automatically fetch either the corresponding tag or media, whichever was not shared, in order to complete the pair, or derive a tag from the media if it was already suitably tagged with meta-data.

The tagset produced on registration allows possession-counting (and potentially, usage-counting), which can be returned to the media servers which are distributed throughout the HP2P network.

Users may trade tags or songs alone, and, given one, the system can find the other on the network. While users can still trade foreign (non compatible) media through the P2P network, they can not use the HP2P systems for searching and classifying. In this sense, this system is very similar to Music2Share [13].

For Mobile-phone systems, using each mobile base transmitter as a media server can be a particularly efficient way to deliver media - bearing in mind the bandwidth/ quality trade-off inherent in a mobile environment. Mobile users would then exchange tags very quickly and the media can arrive automatically at a network-friendly rate.

Tagsets can themselves be shared (in which case, they devolve to playlists) - allowing personalised music selection of new music - or created from a preliminary song-feature selection process. In this case, a new customised tagset can be regenerated periodically, and then distributed from a central repository, implementing a form of personally customised internet radio. This latter method nicely addresses one of the perennial questions in recommender systems based on usage - how to incorporate new music.

Statistics on the customised tagsets would also be of great value to music distributors. They would buy such data from the server hosts (typically ISP's).

3 Conclusion

In this short paper, we briefly describe a HP2P system for sharing music where the user is able to share music freely, but where the marketing value of their ever-changing song collection is used to pay for it.

We also comment on how such as a system can interoperate with a fee-based P2P song-purchase system, to allow the user to easily purchase a higher quality rendition of the song initially obtained freely.

3.1 Acknowledgement

I would like to thank all the people in the User Needs Group of the Smart-Internet Cooperative Research Centre for their support.

References

1. Adar, E., Huberman, A., Free riding on gnutella. First Monday, **5** (2000) http://firstmonday.org/issues/issue5_10/adar/index.html. Accessed 9 May 2005.
2. Brilliant Digital Entertainment, Joltid, Altnet website http://www.altnet.com
3. Amazon Inc. any music search result on the Amazon website http://www.amazon.com
4. Apple Computer Inc, itunes Web Site, http://www.apple.com/itunes/
5. Collab.net Open Source Project, Project JXTA: A P2P Search tool, http://search.jxta.org/project/www/background.html, Accessed 9 May 2005.
6. MetaMachine Inc., eDonkey website, http://www.edonkey.com
7. Garces-Erice, L., Biersack, E., Felber, P., Ross, K., Urvoy-Keller, G., Hierarchical peer-to-peer systems. ACM/IFIP International Conference on Parallel and Distributed Computing (Euro-Par), Klagenfurt, Austria (2003) 1230–1239
8. Gnutella Open Source Project, Gnutella website http://www.gnutella.com, Accessed 9 May 2005.
9. Google Inc., Google website http://www.google.com/, Accessed 9 May 2005.
10. MarketWire Inc., Philips and Gracenote Launch Gracenote Mobile(SM) – First Global Music Recognition and Content Delivery Service for Mobile Phones, http://www.marketwire.com/mw/release_html_b1?release_id=61431, Article dated 9 January 2004
11. Hummel, T, Strømme Ø, La Salle, R (2003), Earning a Living among Peers the Quest for viable P2P Revenue Models, Proceedings of the 36th Hawaii International Conference on System Sciences, Hawaii (2003), p219
12. Sharman Networks, KaZaa website, http://www.kazaa.com
13. Kalker, T, Epema, D, Hartel, P, Lagendijk, R, van Steen, M, Music2Share: Copyright Compliant Music Sharing on P2P Systems, Proceedings of the IEEE, 92:6, (2004)
14. Microsoft Inc, MSN Music, http://music.msn.com
15. Musicmatch Inc, Musicmatch, http://www.musicmatch.com
16. Napster LLC., http://www.napster.com
17. Search Tools Consulting (Avi Rappoport: Principal Consultant), Peer to Peer Searching, http://www.searchtools.com/info/peer-to-peer.html, Article dated 23 May 2002
18. Walmart Inc, Walmart Music Services, http://www.walmart.com/music_downloads/introToServices.do, Accessed 9 May 2005.
19. Xu, C., Wu, J.D. Feng, D., Content-Based Digital Watermarking for Compressed Audio, Proceedings of RIAO2000 Content-Based Multimedia Information Access, pp.390-402, Paris, France, (2000).