Available online at www.sciencedirect.com

ScienceDirect

Procedia Computer Science 52 (2015) 350 – 357

Procedia
Computer Science

The 6th International Conference on Ambient Systems, Networks and Technologies (ANT 2015)

Effects of knowledge base quality on peer-to-peer information propagation

Michael Gibson[a], Wamberto W. Vasconcelos[a]

[a]*Department of Computing Science, University of Aberdeen, Aberdeen, AB24 3UE, United Kingdom*

Abstract

Peer-to-peer (P2P) networks have become popular through the use of file-sharing over a decentralised network. They rely on keyword searching to allow peers to exchange files. However, for networks handling real-time information, such as those from sensor networks or computer games, information cannot easily be categorised by keywords. We developed a knowledge-based P2P mechanism in multiplayer games, based on roles individuals play and their information needs for game playing. In this paper, we show how we evaluated this mechanism when knowledge quality degrades.

Peer-review under responsibility of the Conference Program Chairs
Keywords: Agents; Computer games; Information propagation; Knowledge-base mechanism; Peer-to-peer

1. Introduction

P2P has gained momentum by allowing files to be shared over a decentralised network. This was possible by assigning unique identities (IDs) to each shareable file as well as for each peer within a network so that certain peers will be responsible for holding some of the files and propagate queries for files to the peer with the requested file. Techniques have been developed to assist in query propagation, including keyword-searching so that an identical ID for the requested file does not have to be known in the first place and improving search pattern from 'blind' *flooded request* to neighbour selecting *document routing* [1]. Other improvements include *structuring* a P2P network to organise what peers share among themselves [2] and sharing other types of information, such as from databases [3].

Although *keyword searching* is useful for 'static' forms of information, such as files, this method is not suitable for *real-time* information. Real-time information is produced and consumed in narrower time-frames, and have brief useful lives before becoming out-of-date and useless. Since this type of information has a short useful life, it is unrealistic to identify the information so that it can be queried by peers through keyword searching.

We, therefore, designed a knowledge-based mechanism which uses elements from the document routing mechanism, but can be used for real-time information propagation/exchange, through the use of *roles*. Instead of assigning keywords and IDs to information, we assign roles to peers which *produce* and *consume* types of information. This

E-mail address: {michael.gibson, w.w.vasconcelos}@abdn.ac.uk

Peer-review under responsibility of the Conference Program Chairs
doi:10.1016/j.procs.2015.05.106

means instead of peers querying for information itself, peers will query other peers based on the role which is able to produce information for the querying peer. Another advantage of our mechanism is that peers can *volunteer* information based on what a peer will *consume* according to its role. This is similar to *publish/subscribe*[4], where if a peer knows what another peer will always want, the volunteering peer will automatically give information to the consuming peer without the need for extra query messages.

To test our mechanism, we have developed an experiment to show how the *accuracy* of knowledge can affect the performance of our mechanism and, hence, overall performance of a P2P network. Since we focus on real-time information, we use a *computer game* as the domain of our experiment. The game is a simple resource collecting game, where each player (a peer) has to collect as many resources as possible. Using our role system, we limit players to collecting only one type of resource per role and players must coordinate with each other to collect other types of resources. By changing the accuracy of knowledge a player acquires during a game, we can see how degrading knowledge can negatively affect the overall performance of peers playing the game.

This paper is organised as follows. Section 2 cover related work in sharing knowledge over P2P networks. Section 3 explains how our mechanism works. Section 4 describes our experiment to show how knowledge accuracy affects peer performance. Section 5 looks into how the experiment was evaluated and the results acquired. Section 6 concludes the paper with a reflection on our work and how it can be improved upon.

2. Related work

Typical mechanisms involve selecting the best neighbouring peer to contact in order to request of provide information. The exchange can be a query for a resource, for example, a file, a resource itself or about the network itself which can be used to indicate whether peers are joining or leaving the network. How the best peer is selected depends on the relationship between the information to be propagated and the peers themselves. Examples of such mechanisms include *Freenet*[5], which uses *document routing* to store files with generated IDs on peers with similar IDs in an *unstructured network*[6], and *Chord*[2], which uses *distributed hash-tables*[7] to partition peers in a *structured network*[8] and be responsible for resources that belong to specific partitions. These mechanisms are similar in the sense that each resource and peer can be uniquely identified and their identifications are related to each other.

One of the fundamental issues in P2P networks is how to share knowledge among peers, which is what we are interested in. A growing field in researching this issue is utilising *knowledge bases* to share knowledge over P2P networks. Frameworks, such as *Edutella*[9] and *HELIOS*[10], make use of how knowledge can be organised on each peer to assist query creation and propagation. *Distributed Knowledge Management*[11,12] focuses on how smaller communities of peers in a P2P network contribute to knowledge creation as well as sharing with other communities. With *social networking* becoming an important part of general life, research is also looking into how *virtual communities* can share knowledge among themselves[13,14]. The problem with focusing on small communities within a large community is the phenomena of *small-worlds*[15], where peers will only tend to interact with peers it has interacted with before. To ensure the whole network is up-to-date, it is important for all peers to attempt contact with distant peers.

Since our chosen domain focusses on computer games, we also look into existing work involving games and P2P networks. One of the most important issues in running any multiplayer game is to ensure that all players have the same *game state*. This means all players operate within the same state of the game and any changes made, for example, player location, must be reflected on all players' view of the game. One way to ensure players (peers) are kept up-to-date is to use *interest management*[16,17], which prioritises player updates based on interest factors, such as close player positions in-game and visibly looking at events in-game. Players which exhibit interest in each other will receive more updates to ensure better gameplay, whereas non-interested player will receive fewer updates as their actions will not have as much affect. Another way is to partition game areas through *zoning*[18,19], which makes certain peers responsible for handling all events within a particular area of a game. Interest management and zoning are only useful mechanisms within games (and even to certain types of games). Although they share some similarities with other approaches, such as distributed hash-tabling for zoning, these techniques cannot be easily adapted to other domains. Although we use a game for our experiment domain, our mechanism can be adapted to other domains.