Peer-to-Peer Information Access & Retrieval

S. Johnstone — P. Contreras — F. Murtagh — P. Sage

School of Computer Science
Queen's University Belfast
Belfast
BT7 1NN
Northern Ireland

s.johnstone@qub.ac.uk
p.contreras@qub.ac.uk
f.murtagh@qub.ac.uk
p.sage@qub.ac.uk

ABSTRACT: *In this article the peer-to-peer paradigm is discussed as a means for the provision of complex information access and retrieval systems. A review of peer-to-peer technologies precedes a discussion of their potential to provide the basis from which a rich information service may be built. The benefits of such a system and the problems inherent to peer-to-peer approaches are discussed. The peer-to-peer development platform, JXTA, and its use in the implementation of an image database retrieval system is highlighted. To provide richer information services a system should be aware of its context. Context awareness in sensory information systems is used as a means to model behaviour that will lead to the development of other context aware information systems. To demonstrate the potential of a context aware peer-to-peer system a case study of an active camera network is included.*

KEYWORDS: *peer-to-peer, P2P, context awareness, information systems*

1. Introduction

The area of peer-to-peer (P2P) computing has, in recent times, become a topic of interest, mainly through publicity generated from such file sharing systems as Napster and the subsequent legal cases. The scale that Napster achieved helped to highlight the potential of P2P systems for the development of large-scale information systems.

In order to promote the research and development of P2P applications various efforts have been made to create a common P2P infrastructure. Most notable of these is Sun's JXTA platform that provides developers with the basic peer operations to develop P2P applications.

The nature of the P2P concept suggests that in order for systems to be totally decentralised there must be a certain level of autonomy among the peers. However, P2P work has, thus far, failed to produce any systems that are truly autonomous. Indeed many of the systems, which are mostly file-sharing systems, exhibit very little or no autonomy and are controlled mostly by the user. Systems that have achieved good levels of scale with efficiency have done so through the adoption of a hybrid approach, where there exists a certain level of centralised control. With the presence of some level of central control there will inevitably be a potential for bottlenecks and one of P2P's strengths of fault tolerance will be compromised due to the possibility of failure of the central service.

In order to adequately show the strengths of P2P as an option for large-scale information systems the issue of peer autonomy needs to be studied. Without a reliance on a central service and with the ability to run autonomously the potential for scalability would make such systems an attractive alternative to more traditional client-server approaches. However, autonomous P2P networks with no central authority could pose some security and control problems for certain applications and we do not attempt to suggest that this approach is the best option for all types of information systems.

In this article we describe the prototyping of a simple P2P image sharing application. The platform used was Sun's JXTA and a brief description of this platform and the motivation behind the use of it is given. An overview of the characteristics of P2P networks is also given.

We also highlight the area of ubiquitous computing in which P2P may help achieve some success. In sensory networks, energy restrictions and fault tolerance requirements provide an excellent motivation for a P2P approach. The limited capability of the nodes in these networks also presents a simpler behavioural model. An example application is currently being implemented by us and is described.

2.P2P Systems

At a very basic level a pure P2P system is one in which each entity or peer has the same or similar capabilities as all the other entities in the system. A peer can be thought of as something that can provide some resources but also offer some in return. This immediately sets a P2P system aside from the traditional client-server model where each node in the system would rely heavily upon a central server, whether it be for communication, organisation, data exchange, etc.

If peers are defined to have equal status and are not completely controlled by each other or by some central authority then a level of autonomy becomes apparent in P2P systems. Peers rely on each other to forward requests to other peers for resources or data and each peer will rely heavily upon the behaviour of other peers. Due to their autonomy and their reliance on other peers they cannot fully trust each other. Issues such as peer transience and scale become very important factors to be considered.

P2P networks exhibit many different characteristics that have an impact on the effectiveness of their systems and applications. Some of these characteristics are described below.

2.1. P2P Characteristics

2.1.1.Decentralisation

One of the key characteristics (Dejan *et al.,* 2002) of a P2P system is its obvious move away from the traditional client-server model. While the presence of a central server has the advantage of having a global view of the system allowing issues such as access rights to be dealt with, this traditional architecture can suffer from bottlenecks and unused resources. In a decentralised P2P system if each node has equal status and does not rely on any central service, issues such as organisation become a problem. Such systems must be able to self organise thus giving the peers a level of autonomy resulting in complex models for pure P2P systems.

Some of the problems have been dealt with by the development of 'hybrid' systems where, while the system as a whole is distributed, a certain level of central control is present. For example, the *Napster* system keeps a centralised directory of file locations but the peers directly communicate with each other to download the files. The presence of any central requirement however will have negative effects in terms of scalability as the problem of bottlenecks, seen with client-server models, begins to become a factor again.

2.1.2.Scalability

As mentioned above, the level of scalability of a P2P system can exhibit depends a lot on the level of centralised operations required. A somewhat hybrid system can, however, still achieve high levels of scalability by keeping the levels of central communication to a minimum. Napster, for example, tackles the scalability problem by having the actual downloading of music files performed in a pure peer-to-peer manner and therefore keeping the communication to the central servers at a minimum. At its peak Napster achieved a user level of around 6 million users.

In pure P2P solutions such as the *Freenet* (Freenet 2003) file sharing application scalability is also a limiting factor. In these systems peers can search for and retrieve files. Each peer sends requests to all known peers that in turn forward the request to other peers. This system produces a lot of message traffic and so has scale limits. Current work on this system is based on the development of a more efficient searching mechanism where queries are only sent to those peers most likely to respond.

2.1.3.Self-Organisation

The unpredictable increase in scale, whether it be an increase in load or an increase in the number of peers, means that any sort of central organisation method is very difficult. This has produced a requirement that P2P systems have the capability to self organise. Any increase in scale will, most likely, result in an increase in faults and the increase in peer transience. For a fluctuating system such as this the ability to self-organise is of high importance.

2.1.4.Performance

Obviously performance is important in P2P systems as they aim to be more efficient that traditional systems. Many of the limitations currently put on performance come from the number of messages being sent and therefore bandwidth requirements.

The architecture of the P2P topology employed by the system can play a vital role in the performance of the P2P network. Aside from the pure P2P topology where every peer will have completely identical status in the network an attempt to improve the performance of P2P networks has resulted in the use of the Super Peer topology where it could be argued that while all peers may be equal, some are more equal than others. A discussion on the super peer topology is included in this article.

2.1.5.Interoperability

Although a number of P2P architectures and applications have been produced none of these can operate with each other. A FreeNet peer cannot communicate with a Gnutella (Gnutella 2003) peer for example. This lack of common infrastructure is where P2P currently lags behind other distributed computing solutions such as grid technologies (Foster *et al.,* 2003). Due to the current relatively simple P2P applications, standard infrastructure and tools may not be required, but as applications become more complex the benefits from common standards increase. While the scale, in terms of numbers of nodes, in current grid applications is relatively small compared to the millions of users of some P2P systems, grid developers can tackle the problems of scale and infrastructure in parallel as the infrastructure already exists. In contrast, where P2P systems have achieved high scalability the lack of any common infrastructure will ultimately impose a limit on the scale due to lack of interoperability.

There are efforts, however, to introduce common platforms for the development of P2P systems namely, Sun's JXTA (JXTA 2003) and Microsoft's .NET (Microsoft 2003). There are also academic based platforms such as Berkeley's BOINC (BOINC 2003) and Paris XI University's XtremWeb (XtremWeb 2003). While the JXTA platform has attracted a lot of interest it has yet to really achieve widespread acceptance as a standard.

An overview of the JXTA platform is given in the following section. This platform was selected for our initial investigation into P2P application development due to the relatively high levels of interest and the availability of the source code. The JXTA platform is by no means a finished system and in fact parts of its implementation are updated on a daily basis.

2. 2.JXTA P2P Platform

P2P computing is a network model where possibly any node (peer) may be interconnected. A peer may also operate as a consumer or provider, unlike the traditional client/server architecture.

JXTA (JXTA 2003) is an open source project (created by Sun Microsystems) that defines the following set of six language independent protocols (Protocols 2002) to provide a P2P infrastructure:

- *Peer Discovery Protocol (PDP)* Peer/Resource discovery
- *Peer Resolver Protocol (PRP)* Generic query service

- *Peer Information Protocol (PIP)* Net monitoring
- *Pipe Binding Protocol (PBP)* Addressable messaging
- *Rendezvous Protocol (RVP)* Propagation service
- *Endpoint Routing Protocol (ERP)* Routing

These protocols establish a ubiquitous, secure and pervasive virtual network (Traversat et al., 2002) on top of internet and non-IP networks, allowing peers to directly interact and be organized independently of their network locations, even behind firewalls and NATs (network address translation).

JXTA has been designed with the objective of making it ubiquitous. Any device including mobile phones, PDAs, pagers, electronic sensors through to desktop computers and servers, has the potential of being JXTA enabled. In order to maximise its platform independence and interoperability, widely accepted standards like XML form the basis of JXTA communications. Security (Security 2001) is considered through the use of Transport Layer Security (TLS), Digital Certificate and Certificate Authorities.

JXTA protocols may be used to build different types of application, e.g. instant messaging, file sharing, data search, remote sensor communication, agent based systems (Babaoğlu *et al*., 2002). It can also be used to implement a wide range of services. Thus, this technology provides the ability to integrate heterogeneous information sources in a decentralized, self-organized and secure way.

2.3.JXTA Architecture

Figure 1 shows an abstraction of the JXTA platform architecture (Wilson 2002).

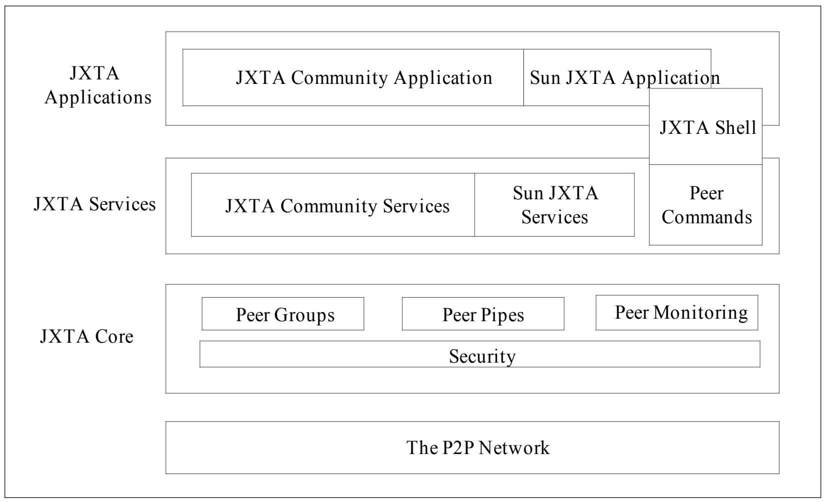

Figure 1. *JXTA Architecture* (Wilson 2002)

2.3.1.Core Layer

This layer contains elements that are essential to a P2P solution. The elements of the core level are listed below:

- Peers
- Peer Groups
- Network transport (pipes, endpoints messages)
- Advertisements
- Entity naming (identifiers e.g. peer IDs)
- JXTA protocols
- Security and authentication

2.3.2.Services Layer

The services layer provides network services that are considered to be desirable but not totally necessary to all P2P solutions. Such services provided include:

- Searching for resources on a peer
- Sharing documents from a peer
- Performing peer authentication

An example of a JXTA service is the Content Management System (CMS 2002) that allows the sharing and retrieval of content within a peer group

2.3.3.Applications Layer

The dividing line between the services layer and the application layer can be hard to distinguish. Usually anything in this layer builds upon the functionality provided by the service layer. Instant messaging is an example of something that lies in the application layer.

JXTA uses asynchronous uni-directional pipe communication channels known as pipes for sending and receiving messages. The presence of peers, peer groups and pipes are all advertised through the use of XML advertisements. Sun has provided an open source Java implementation of the JXTA protocols as well as some core services to provide extra functionality. Examples include CMS and JXTA Search that provides a more efficient searching mechanism than a traditional P2P search that sends requests to all known peers hence creating excessive network traffic.

To serve as both a proof of concept and to make some evaluation of JXTA as a development platform a series of additions/modifications to the JXTA platform, CMS and InstantP2P (InstantP2P 2002) user interface were made in order to set up a Simple Image Service that would supply a user with some image retrieval and processing functionality in a P2P manner. InstantP2P (aka. myJXTA) is an application freely available from the JXTA site that gives the user a GUI (graphical user interface) giving access to peer discovery, group membership, file sharing and chat functionality through a window-based interface.

3.Simple P2P Image Service

A service was implemented where a user can search for an image in a P2P environment. To give more flexibility to this scenario, once the image has been located, different wavelet transform and multi-resolution (MR 2002) analyses including decompression of compressed image files can be applied invoking MR software.

In this context, the Java binding (Builds 2002) of JXTA is used to create a virtual P2P network and the JXTA Content Management System is used to transfer the digital content. In addition, some extensions to the CMS core libraries are implemented to announce particular services.

Figure 2 shows how using this technology, the user is provided with a flexible application in order to access and analyse images. First, a provider advertises some image files by name and/or a free text content description. The multi-resolution service for this particular image is included in this *content* advertisement. When a user searches for a particular image a JXTA query is propagated. When the content advertisement is located, the user can apply a multi-resolution analysis if desired, the image is compressed and sent to the requesting peer. When the image is received it is decompressed and displayed to the user.

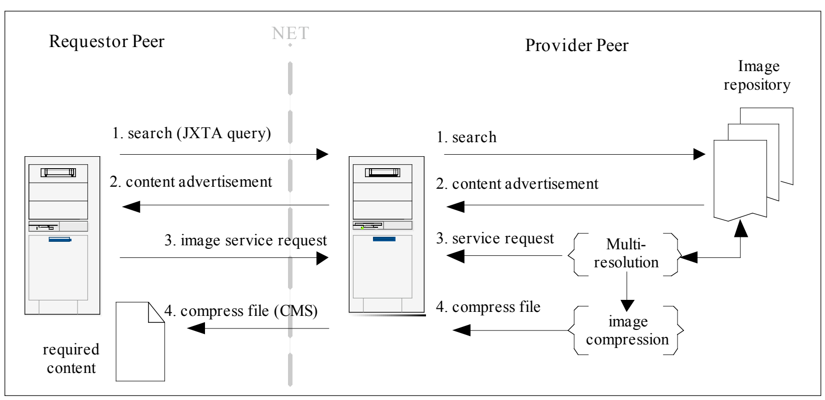

Figure 2.*Simple P2P Image Service*

3.1. Distributed P2P Database Access

To incorporate flexibility database searching is added to the model. In our particular approach the MySql database system and JDBC drivers are used to connect to and retrieve information from the database. In addition, multi-resolution processes and compression-decompression, as described above, may also be applied.

Figure 3 shows the process where a user searches for a particular multimedia content. The query is propagated through the peer network group and then received by the provider peer. The provider looks into the database and if this content exists, the provider creates a content advertisement with the result, which is returned to the user. Once this is done, the user can select a particular content advertisement and use this in order to retrieve the file.

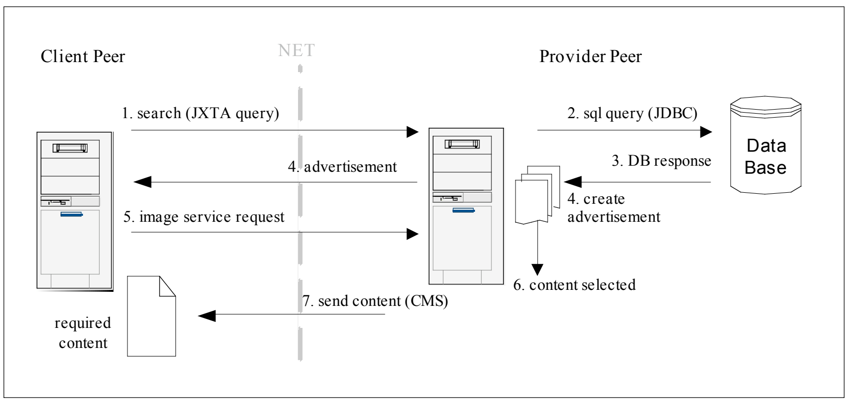

Figure 3. Retrieval Process in a JXTA Database Context

3.2.P2P Distributed Image Repository

A tailored architecture for a wide range of applications could be created, taking into consideration the above examples. We have taken an astronomy example and implemented a P2P distributed image repository. JXTA provides a transport layer for collaborative and distributed work on image data. The image access infrastructure uses standards like "Simple Image Prototype Specification" (SIAPS 2002) to provide services specification and VOTable (Williams *et al*., 2002) is used as the communication standard. We can access and query databases through the peer as well as local repositories.

Figure 4 shows how the P2P distributed image repository works. The figure has been split into three sections to clarify the architecture. Section *A* shows a user performing a search for available image services within their current peer group. This query is marked ‘1’. The response received by the user is a list of the currently available image services, marked with ‘2’.

In section *B* the user has chosen a specific service to which a request using the SIAPS specification is sent. The response to this query is in the form of a VOTable, which includes a list of image names and their respective properties. The next step is described in section *C*. After the user has selected the files that they wish to retrieve, a retrieval request is sent to the provider and the requested files may be sent to the requesting peer. For the image retrieval process the JXTA CMS libraries are used.

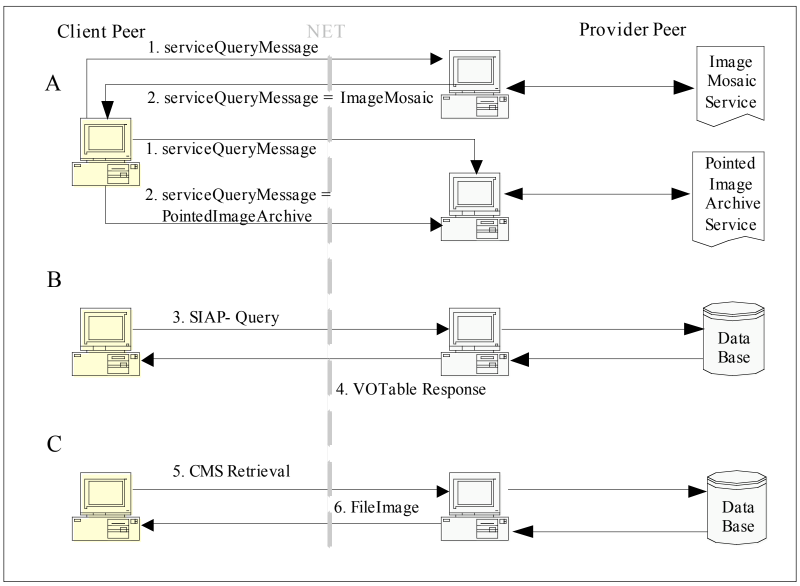

Figure 4.*P2P Distributed Image Repository*

4.Context Aware Information Providing

The implemented work described in the previous sections highlights the potential autonomy that individual nodes in a P2P information network may achieve. Any peer in these simple applications implemented can be either a consumer or a provider, thus reducing their reliance on any central authority or service and therefore improving fault tolerance of such systems by removing a single point of failure. This level of autonomy may be further increased if the peers can become aware of both their own and their users' physical or semantic position and adapt their behaviour accordingly while being minimally intrusive (Englmeir 2003). Such context awareness would allow peers within a network to be aware of their users' and/or applications' context and use this knowledge to make decisions regarding peer-to-peer communication and self-organisation.

To produce or even model such a system would be a complex task. For example a user's context could consist of many attributes such as current position, personal history, analysis of behavioural patterns and interest profiles (Englmeir 2003). We believe that to best investigate the area of such P2P context aware information systems it is beneficial to look at some simple application areas where basic application/user models can be produced and limited peer behaviour is possible, thus reducing the overall size of the system model and producing a good starting point from which more complex models may be based. It should be possible from such a platform to investigate some of the issues that would need to be addressed such as P2P computing issues like fault tolerance and the machine-learning aspects concerned with context awareness.

5.Sensory Networks

An area that could produce simpler peer behavioural models is that of sensory networks. The nature of the nodes that make up these networks (the sensors) and their limitations means that their capabilities are reduced and so to model these as peers is an easier task. CEO of Graviton (a company specialising in wireless sensor networks), Solomon Trujillo, likened future sensor networks to a digital human nervous system - '*Consider the human body. Most of what goes on in the human body is done automatically and not at a conscious level. For example, our heart beats, our pupils dilate and constrict, when hot our sweat glands go to work creating the mechanism that causes the evaporation of sweat from our skin that lowers body temperature. If we had to consciously control all of these bodily functions our conscious state would rapidly grind to a halt'* (Trujillo 2001).

Trujillo's uses this metaphor to explain his vision of the future of sensor networks. He believes that machines, through the use of sensory information, should be able to process and exchange information with each other, sense changes in their environment and adjust to those changes. If we consider the aims stated above in the context of information provision and access it can be seen that the goals are convergent and give further justification for the initial prototyping of sensory network applications to build a solid foundation for more complex distributed context aware information systems.

6.Case Study: Active Camera Network

The motivation behind this application area is two-fold. Firstly in an active camera network where directed attention of cameras is a required level of functionality, communication between cameras may be considered as a P2P style of communication where, ideally, the cameras will act autonomously (in the absence of some central authority) and communicate with others in the network as and when they deem it appropriate to do so. In this application area the active camera network may be viewed as a peer group where the peers are the cameras themselves or processing nodes directly controlling the cameras. The peer behaviour in this application is limited and so will produce a simpler data model and therefore provide a good basic starting point for the work.

In this scenario a network of surveillance cameras form the source of sensory information. The goal for this type of system is to be able to achieve directed attention of the cameras. For example, if a camera picks up a person walking along a corridor it may wish to track the person's movement. Depending on the direction of movement the camera may wish to alert other cameras that have the ability to track the person. The work of interest here is the major parts that would make up this system and how to model the mapping between them. The system would be organised in a P2P fashion where either the cameras, if equipped with processing capability, could represent the peers or a set of external peers could control the cameras directly. A system like this can be abstracted into 4 parts:

1. Application: High-level objectives for the system as a whole to achieve. An example for this scenario could be *'secure entrance hall and adjoining corridors'*.
2. Sensory Context Analysis: At this level the system performs analysis such as some image tracking in order for the camera to become 'context aware'.
3. Event Notification: Based on the context analysis an event may be triggered in order to achieve the application level requirements. At this stage a single camera tracks the object or it may notify other neighbouring cameras that there is an object of interest.
4. Middleware: The layer that maps the sensory events onto application events.

The two major requirements for such a system for which we would be concerned are as follows (Mohan *et al*., 2000):

1. Adaptive: Systems should be able to adapt to changes in environment. This may include cameras or other sensory devices being added to or removed from the network.
2. Active: Systems should be able to direct attention to interesting events

6.1.Adaptive

The P2P approach should help to achieve the requirement that the network be adaptive in so much that a P2P network will not fail if peers drop out of the network and also easily accept the arrival of new peers. This scenario creates a more complex problem in that a new peer will be actively participating in the network and will need to make its presence and capabilities known. This opens up some more problems:

- What the network should do if a node fails?
- What peers will a new peer need to communicate with in order to fit into the network?
- What data will a new peer need to exchange when joining the network?

These questions demonstrate the need for autonomy achieved through awareness. For the system to be able to self organise peers will have to have some way of discovering their own physical and semantic context. They will need to know about the capabilities of the resources around them.

6.2.Active:

Realising this requirement would involve modelling the context in which the sensory data is collected and how data from the environment context can be used in order to meet the application level context. The key issues here are:

- How do the sensors discover their context based on sensed data?
- What combination of data constitutes an interesting event?
- Depending on events what communication is required between peers?
- How to link sensed events to application level events?

6.3.Architecture

The fact that an application layer exists suggests that at some point the system will have to have a single point of origin i.e. where the application requirements come from. This would suggest that a hierarchy might be formed. Each peer will have to have at least some limited knowledge of the application requirements such as events that must raised. It would be inefficient in terms of a scalable network for the application source to broadcast its requirements to each and every peer. This would motivate the adoption of a super peer style of network topology shown in Figure 5.

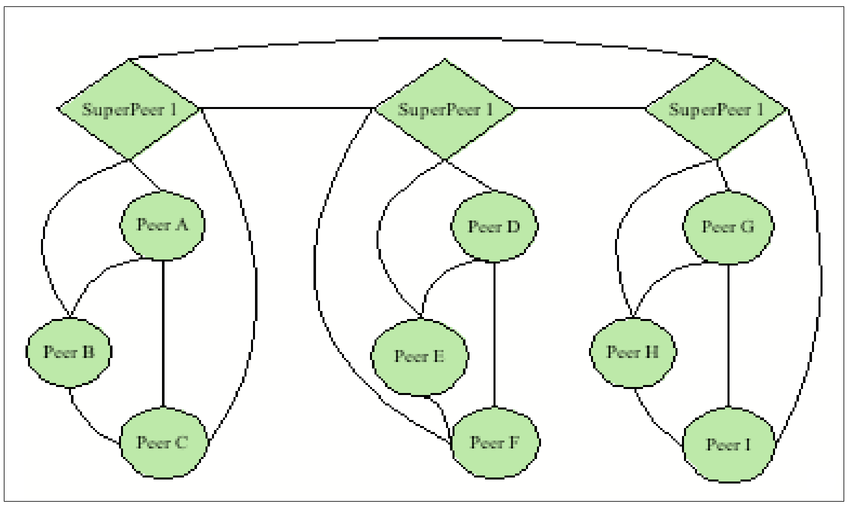

Figure 5. *The Super Peer Topology*

When considering this type of topology the following points should be considered:

- **Election of Super Peers** – We propose that these decisions could be based on results of a context analysis process in the system. Traditionally such an election of super peers may have been simply a manual selection by the users of the system or automated approaches based on considerations such as hop times, available bandwidth etc. such as those described in P2PRise (Babaoğlu *et al*., 2002). The proposed context based election should produce P2P networks organised to best suit their particular applications.

- **Points of failure** – With the adoption of the super peer topology the problem of single points of failure must be considered. For example the failure of a super peer could leave peers that were directly connected to it isolated from the rest of the peer group. A major question to be considered is, can failure be efficiently and quickly detected hence resulting in the automatic election of a new super peer to resolve the failure problem?
- **Low-level peer/application context Awareness –** If organisation decisions are to be made based on a context analysis phase then how this analysis is performed is of utmost importance. A rule based AI approach may be appropriate.

This proposed super peer structure directly motivates the use of JXTA. JXTA 1.0 used a completely pure P2P approach but latest versions have adopted a super peer topology. In JXTA the super peers are known as Rendezvous Peers. Any peer can be a rendezvous peer. JXTA also allows completely customisable messaging with defined interfaces for developers to design any sort of message to be sent in XML to other peers.

With the low level peer architecture considered attention must be directed to an overall high-level system view. The peers and the application must be able to provide the correct information to the system, or the system must be able to extract the correct information from the peers and the application. This suggests that for the system to function correctly the application and the peers must be able to meet in the middle. Already the P2P platform can be seen as a layer in the system between the application and the peers. It may be useful therefore to augment this middleware layer with machine-learning techniques necessary to make the system context aware. The aim of this module would be to produce a set of data relating to application requirements and low-level context analysis based on a set of rules/questions that it would ask both the low levels and the application level. A layered overview of the system may be seen in Figure 6.

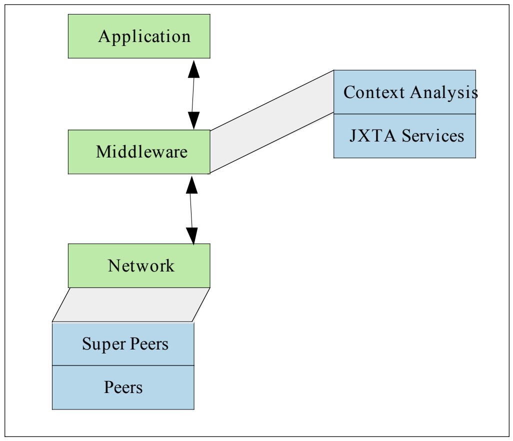

Figure 6.*System Layer Overview*

6.4.The Active Camera Network Prototype

In order to begin the implementation of the prototype a simple scenario is considered. The application is one of building security. Each camera or peer will be located in a particular part of the building. Depending on each camera's position there will be a number of other cameras in the network that will be interested in detected objects moving in their direction. When a particular camera detects an object of interest it will alert the cameras in its neighborhood that an object may enter their area from a given direction. The alerted cameras should then direct their attention to the given direction.

Implementation of a prototype system is currently underway. The work to date has concentrated on peer level organization, communication and context. In its current stage knowledge of the necessary application context is assumed. The JXTA platform provides us with the necessary peer discovery and communication protocols from which this application can build. Before discussing the current state of the prototype first consider the application's current structure shown in Figure 7.

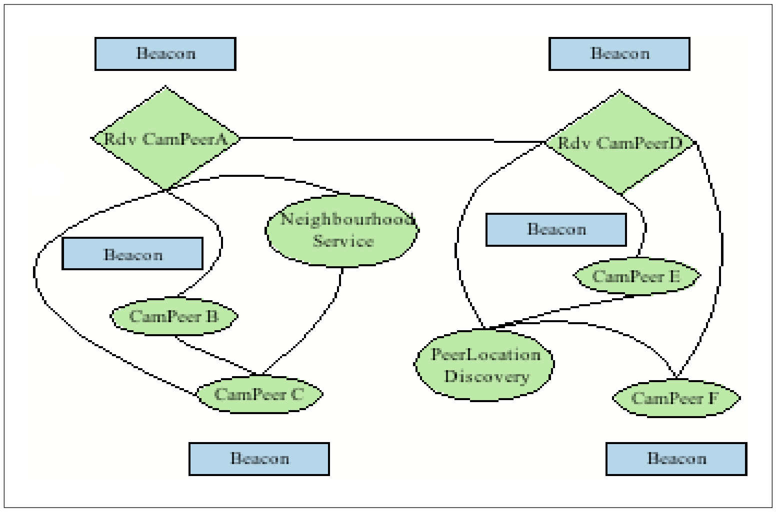

Figure 7. *Prototype Structure*

The JXTA platform segments the network space through the use of peer groups. When a peer is instantiated it automatically becomes a member of the netPeerGroup. The netPeerGroup is seen as the parent group for all other groups. This prototype is contained within one peer group named ActiveCam. The prototype has been designed so that when a peer is instantiated it automatically searches for connection to a super peer known as Rendezvous Peers in JXTA. In JXTA if a rendezvous peer is not present it is possible for a peer to dynamically become a rendezvous peer. Note that in the diagram it can be seen that the rendezvous peers are also Camera Peers. This is just to illustrate that any peer can be a rendezvous peer. When connected to a rendezvous peer a camera peer will perform a search using the JXTA discovery service to try and locate the peer group ActiveCam. If found the camera peer will join this group and if not located it will create the group. When in the group the Camera Peer can now communicate with the other peers within the group.

6.5.Peer Location Context

In order to begin the investigation into a context aware P2P system, location awareness was chosen as the initial context type for two reasons. Firstly it is a relatively simple context to analyze in comparison to other sensed data and secondly due to the availability of existing work in the area.

In this prototype it is necessary for the camera peers to first of all know their absolute physical location within their environment and secondly to know which peers are in their local neighborhood. The cameras must be aware of both their physical location and their semantic location. In the case of physical awareness, in an outdoor setting, such technologies as GPS could be used in order for a peer to find its physical location. However this prototype, at present, is concerned with indoor security and so GPS is not an option. For this problem we have looked to HP's Cooltown (Pradhan 2003) project which is largely concerned with location awareness. If we consider that for this type of application the end user, in the event of a security violation, will need to know the location of the source a network location will be meaningless but a location name in the context of the building within which the system operates will carry some meaning to the user. Therefore, when a camera joins the network it must receive some semantic data relating to its location. In Cooltown this is achieved through the use of IR beacons that transmit location context data. When a mobile device comes into the range of one of these beacons it receives data concerning it location. We have used this idea, as can be seen in Figure 7. When a camera peer enters the network it waits until it receives its context data from its local beacon. Along with context data it also receives a set of physical co-ordinates so that its local neighborhood may be calculated.

6.5.PeerLocation Discovery and the Neighborhood Service

In order for a camera peer to be aware of its neighborhood at some point in the whole system a global view of the peer group is required. The acquisition of this global view is the task of both the PeerLocation Discovery and the Neighborhood Service as shown in Figure 7. Viewed abstractly and considering Figure 6 both these modules would be part of the middleware layer.

When a camera peer receives its location data it then locally publishes an XML PeerLocation advertisement. This advertisement contains its received location data along with its uniquely identifiable JXTA peer ID. PeerLocation Discovery is a process whose task it is to keep an up to date table of the current peers and their locations within the group. It does this by periodically, using the JXTA discovery service, searching for these PeerLocation advertisements. In the event of a peer no longer residing in the group it will be removed from the table and new peers will be added.

The PeerLocation Discovery module is intended for direct use by the Neighborhood Service. This has been set up as an actual JXTA service and so a service advertisement is created both locally on the service's host peer and remotely on its rendezvous peer. This can then be easily found by any of the peers within the group. This service receives requests directly from camera peers within the group and responds with a set of neighboring peers for the requestor. Direct communication is achieved through the use of the JXTA PipeService. A JXTA pipe is an abstraction used to send messages directly to other peers. In this case a bi-directional pipe is used. Included in the service advertisement is a unique pipe ID that peers wishing to send a message to the service must use in order to receive their peer neighborhood.

When a camera peer has received its peer neighborhood it is now free to begin its task of object detection and alert its neighbors directly when it decides it is necessary to do so. This is the next stage of the prototype development and will involve a much more complex context analysis phase in order to extract meaningful context from sensed data.

6.6.Peer Communication

During normal operation it is required that the peers can communicate directly with each other. On viewing the super peer topology used in this prototype it would seem that, for example, referring to Figure 7, CamPeerC cannot communicate with CamPeerE without first going through both rendezvous peers. This is true but only to a small extent. It is true that in order to discover the presence of CamPeerE, CamPeerC will need to send a discovery request message to RdvCamPeerA who will in turn forward the request to RdvCamPeerD. RdvCamPeerD will respond with an index to CamPeerE, which will be returned to CamPeerC. Using this CamPeerC can then communicate directly with CamPeerE without further use of the rendezvous peers. This is achieved through the use of the rendezvous peers' SRDI or Shared Resource Distributed Index mechanism where indexes to their known peers are stored.

6.7.Future Development

The immediate future of development will be adding support in the middleware layer to allow the identification of application context. For example in this security application the layout of the building to be secured will have to come from the application level. It will be the task of the middleware to translate this to data that can be used in the location awareness of the peers in the lower level. The application may, for example present its set of requirements to the middleware layer in a fixed format which can then be converted to a form which the peer level can make use of.

Of course, to make this a useable application the peer level context analysis module will have to be augmented to be able to make sense of sensed data such as object detection. It will need to be able to decide if an object is a threat and whether it needs to alert the peers within its neighborhood. Work carried out on the TEA project (Albrecht *et al*., 2001) may prove to be useful with this context awareness level. Here raw sensor data is first abstracted slightly to sets of cues from which contexts may be derived.

Longer-term developments of this work will seek to model more complex peer behavior. For example where a peer is a PC collaborating in a P2P multimedia sharing network both the peer behavioral model and the application model will be more complex. Work in this direction will serve both to provide good research into achievable levels of autonomy in distributed information systems and to provide a convincing justification for the P2P paradigm itself.

References

Albrecht Schmidt, Kristof Van Laerhoven. How to Build Smart Appliances? *IEEE Personal Communications*. August 2001, vol. 8 no. 4, p. 66-71

Aravind Krishnaswamy, University of Waterloo. P2PRISE: A Proposed Framework for a Self Configuring Peer-to-Peer Rendering System. http://www.intellifix.com/Papers (April 2003)

Babaoğlu Ö, Meling H., Montresor A. (2002). Anthill: A Framework for the Development of Agent-Based Peer-to-Peer Systems. *22nd International Conference on Distributed Computing Systems*. Vienna, Austria, July 2002.

BOINC (2003) Berkeley Open Infrastructure for Network Computing. BOINC website. Retrieved September 12 2003 from http://boinc.berkeley.edu/

Builds. (2002). JXTA Project Stable Builds. Retrieved January 25, 2003, from http://download.jxta.org/stablebuilds/index.html

CMS. (2002). Content Management System. Retrieved January 25, 2003, from http://cms.jxta.org/servlets/ProjectHome

Dejan S. Milojici, Vana Kalogeraki, Rjan Lukose, Kiran Nagaraja, Jim Pruyne, Bruno Richard, Sami Rollins, Zhichen Xu. Peer-to-Peer Computing. HP Technical Report. March 8 2002

Englmeier Kurt. Setting the scene for context-aware information providing. *10th International Conference on Human-Computer Interaction*. Crete, Greece, 22-27 June 2003

Freenet. (2003) Freenet website. Retrieved September 12 2003, from http://freenetproject.org/qcms/

Gnutella. (2003) Gnutella website. Retrieved September 12 2003 from http://www.gnutella.com

Ian Foster, Adriana Iamnitchi. On Death, Taxes and the Convergence of Peer-to-Peer and Grid-Computing. *2nd International Workshop on Peer-to-Peer Systems (IPTPS'03)*, February 2003, Berkeley, CA

InstantP2P Project Home. Website retrieved on Spetember 12 2003 http://instantp2p.jxta.org/servlets/ProjectHome

JXTA. (2003). Jxta website. Retrieved January 25, 2003, from http://www.jxta.org

Microsoft .NET. (2003) What is .NET? Retrieved September 12 2003 from http://www.microsoft.com/net/basics/

Mohan M. Trivedi, Ivana Mikic, Shailendra K. Bhonsle. Active Camera Networks and Semantic Event Databases for Intelligent Environments. *IEEE Workshop on Human Modeling, Analysis and Synthesis* (in conjunction with CVPR), Hilton Head, South Carolina, June 2000

MR. (2002). Multi-resolution Image and Data Analysis. Retrieved January 25, 2003, from http://www.multiresolution.com

Pradhan Salil. Semantic Location. Retrieved July 30, 2003, from http://cooltown.hp.com/dev/wpapers/semantic/semantic.asp

Protocols. (2002). JXTA v1.0 Protocols Specification. Retrieved February 1, 2003, from http://spec.jxta.org/v1.0/docbook/JXTAProtocols.html

Security. (2001). Security and Project Jxta. Retrieved January 25, 2003, from http://www.jxta.org/project/www/docs/SecurityJXTA.PDF

SIAPS. (2002). Simple Image Access Protocol Specification . Retrieved January 24, 2003, from http://bill.cacr.caltech.edu/cfdocs/usvo-pubs/files/ACF8DE.pdf

Traversat, B., Abdelaziz, M., Duigou, M., Hugly, J., Pouyoul E. & Yeager, B. (2002). Project Jxta virtual network. Retrieved January 20, 2003, from http://www.jxta.org/project/www/docs/JXTAprotocols_01nov02.pdf

Trujillo Solomon D., Chairman and CEO, Graviton. Address to The Executives Club of Chicago, Chicago, Illinois. May 17 2001. The Future of Sensory Networks

Williams, R., Ochsenbein, F., Davenhall, C., Durand, D., Fernique, P., Giaretta, D., Hanisch, R., McGlynn, T., Szalay, A. & Wicenec, A. (2002). Retrieved January 21, 2003, from http://www.us-vo.org/VOTable/VOTable-1-0.pdf

Wilson Brendon. *JXTA*. , New Riders 1st edition, June 2002, ch. 3, p. 5

XtremWeb, (2003) A Global Computing Experimental Platform. Website. Retrieved on September 12 2003 from http://www.lri.fr/~fedak/XtremWeb/introduction.php3/