

# Providing Presence within P2P systems

Computing Department, Lancaster University

## 1.0 Introduction

The ability to convey presence information is increasingly becoming an important aspect of many systems. In particular, presence has played a role within the areas of CSCW (where it is commonly referred to as awareness), distributed systems (being aware of what services exist) and Grid computing (being aware of which nodes in the grid are available). Examples of presence being used in P2P include instant messenger applications (such as ICQ and AIM) and distributed computation applications (such as [SETI@home](#)). The use of presence brings with it contextual information which in turn can provide advantages within a system. For example, it can assist co-operation by allowing users to see whether or not other users are busy, or be used to optimise distributed computation by allowing a system to be aware of when a node is connected [Palen].

Presence information can be particularly important for large organisations that are globally distributed. A notable problem that is often experienced is the amount of time it takes to resolve issues that involve people from more than one site [Herbsleb]. In such circumstances, presence information could be used to alleviate problems by informing distant colleagues who is available, and when they are available [Godefroid].

Presence, itself, can be loosely considered as information that is made publicly available about a particular user, peer or resource. Furthermore, this release of information is controlled by these respective sources. For example, a user assigns their state (free, busy, etc) and this is broadcasted within a system. It is the user who decides what information they want to make available (so for example, they may state they are off-line just to ensure they are not disturbed).

This document seeks to discuss how presence can be achieved within a P2P system. It first examines the main types of presence before examining what would be required in different types of logical network architectures. Finally, the document ends with a discussion of issues that can affect presence within a system. These include privacy, controlling information and real-time consistency.

## 2.0 Types of Presence

Presence within a P2P system can be split into two main categories, *peer presence* and *abstract presence*.

**Peer presence** represents information that is available about the peers that are on the network. This information typically includes whether or not the peer is currently online, but can also include other information such as the IP address of that peer or perhaps information about its network connectivity.

**Abstract presence** represents information that is available about the entities that utilise the peers, or represents information that is available about the peers'

environment. Types of abstract presence can include users, resources or even the peers' physical environment.

- *User presence* represents information that is available about the user of a peer. This typically focuses on whether or not the user is online, but also can represent information about a users state. For example, this could be whether or not the user is busy, or is away from their computer. There are some issues that need to be considered with user presence, however. For example, although the peer may be connected to the network, the user could have registered themselves as being off-line. Furthermore, it is perfectly possible that multiple users may use a peer.
- *Resource presence* represents information that is available about the resources that are connected to the peer. What constitutes a resource is difficult to fully define, however they do seem to fall into two categories, *internal* and *external* resources. An internal resource can be regarded as being those resources that contribute to the actual system, for example processing power, hard disc space, bandwidth, services, etc. External resources can be regarded as being those resources that are independent of the system and play no role within its function. These, for example, can be MP3's or documents that can be shared. Typically a P2P system supports the communication/utilisation of external resources.
- *Physical presence* represents information about the physical environment of the peer. This could include location information [Want], audio and video information (for example, with the use of web cams) and information about the current environment (for example, what web site is the user currently browsing, what file are they editing).

Although these two categories of presence are quite distinct they are intrinsically linked together as the abstract layer cannot exist without a peer layer. Abstract presence can essentially be considered as an extension of peer presence. However, the relationship between the two does not have to be one to one. This is illustrated in figure 1.

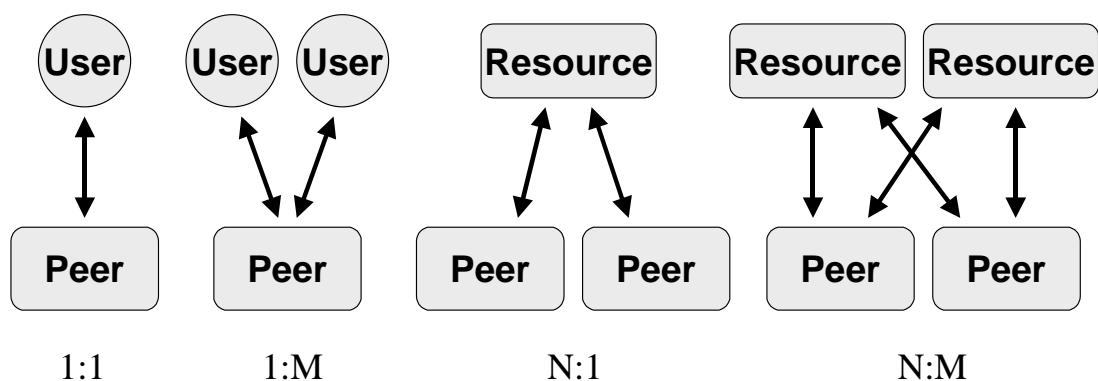


Figure 1 - Possible relationships between the abstract and peer layers

An example 1:1 relationship between peer and abstract layers is that of a single user making use of a single peer, e.g., a user utilising an instant messenger application on their PC.

An example 1:M relationship between peer and abstract layers is where multiple users make use of a single peer, e.g., more than one user utilising an instant messenger application on a shared PC.

An example N:1 relationship between peer and abstract layers is where a resource is controlled by more than one peer, e.g., a hard disc that is shared between two peers.

An example N:M relationship between peer and abstract layers is where a resource is controlled by more than one peer and peers control more than one resource, e.g. peers might share a hard disc and a printer.

### **3.0 Providing presence at the peer and abstract layer**

When considering presence within P2P systems, there are two main aspects that need to be taken into account. Firstly, the deployment of the presence information throughout the system and secondly, reacting to any changes to this information. To an extent, both of these aspects will be affected by the choice of underlying logical architecture used.

Deploying presence information – similar to broadcasting changes (E in the issues raised below) and so therefore will not be discussed.

Reacting to changes to presence information can be further broken down into the following issues within the areas of: identifying when an information change can occur and then as a result of this, ensuring that interested parties have up to date information.

#### **When presence information may change**

- A. *When the entity providing the presence information is connected to a P2P network.* For example, a peer's state may change from off-line to on-line when the application is started
- B. *When the entity providing the presence information is intentionally disconnected from a P2P network.* For example, an application is terminated and so the peer's state changes to off-line.
- C. *When the entity providing the presence information is accidentally disconnected from a P2P network.* For example, if there is a power cut and the peer is accidentally disconnected from the network. In this case it is likely that the peer would not have informed the rest of the network about its change in state. This means that it may be necessary to also provide an alternative mechanism for keeping interested parties up to date with the latest presence information.
- D. *When the entity providing the presence information changes the information that it publishes.* For example, if a user changes their state from busy to free.

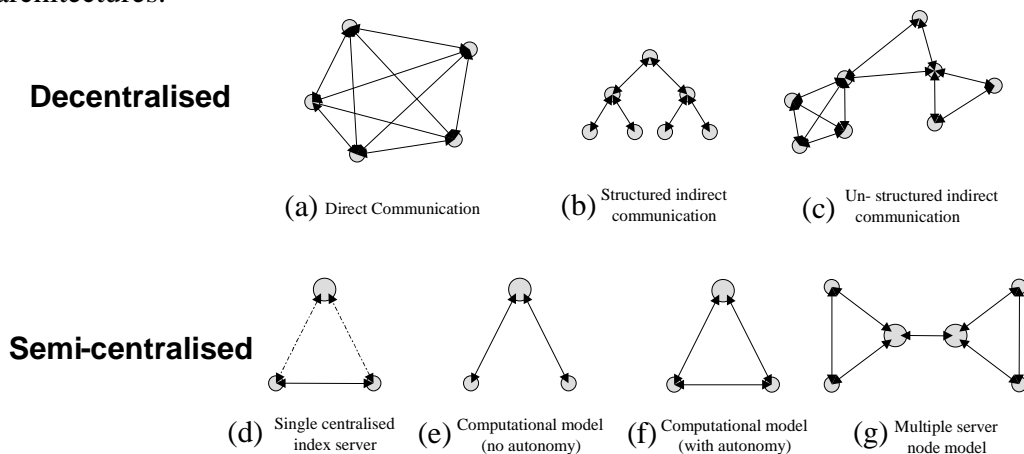
#### **When it may be necessary to ensure that interested parties are kept up to date**

- E. *When an entity connects to a P2P network it needs to be informed about the current presence state of the network.* For example, when an instant messenger application is started, it is informed of which other (relevant) users are on-line. Essentially this represents an entity being given the latest presence information when it connects to the P2P network.

- F. When the presence information being published by an entity changes. For example, if a user changes their state from busy to free. In this case interested parties need to be informed of this change. This can either be achieved by just informing them of the change, or by republishing all the presence information.

## 4.0 Summary of P2P logical network architectures

In this section we summarise the different logical network architectures that a P2P system can be built upon. A more detailed review of the different architectures has been provided in the deliverable Report on the Dependability properties of P2P architectures.



**Figure 2 - P2P Architectures**

*Direct communication* –All nodes are regarded as equal and autonomous. No single node maintains any control over the network. Each node can communicate directly with each other. Each node is aware of each other. As a result of these characteristics, scalability is likely to be an issue.

*Structured indirect communication* - All nodes are regarded as equal and autonomous. No single node maintains any control over the network .It is not necessarily the case that all nodes can communicate directly with one another. Communication could be routed via other nodes. Nodes are connected together in a structured manner (for example, hierarchical, star, ring, etc). A degree of management may be required to ensure the structure persists.

*Unstructured indirect communication* - All nodes are regarded as equal and autonomous. No single node maintains any control over the network. It is not necessarily the case that all nodes can communicate directly with one another. Communication could be routed via other nodes. No structure is forced onto the architecture. The discovery service becomes particularly important

*Single centralised index server* - A single node acts as a lookup for all other nodes within the network. All other nodes are regarded as being equal and autonomous. All nodes can communicate directly with each other, but the index node typically facilitates this. These index nodes are a single point of failure for the architecture.

*Computational model (no autonomy)* - A single node acts as a focal point for all other nodes within the network. The remaining nodes do not possess their own autonomy. All communication is via the server node, if at all. Arguably not a true P2P architecture. The server node is a single point of failure for the architecture.

*Computational model (with autonomy)* - A single node acts as a focal point for all other nodes within the network. The remaining nodes retain a degree of autonomy. Nodes could communicate directly with one another (typically facilitated by the server node). The server node is a single point of failure for the architecture

*Multiple server's model* - Not necessarily the case that only one server node can exist. Allows the possibility of hybrid architectures. For example, server nodes connect together via a direct communication architecture, but collectively act as a single server node within a semi-centralised architecture

## 5.0 Establishing presence support in a P2P architecture

This section focuses on the main types of P2P logical network architecture and suggests ways in which the presence issues identified in section 3 can be satisfied within each architecture.

### 5.1 Providing presence within direct communication decentralised systems

Of all the different types of architecture, those that utilise direct communication between nodes are likely to always provide the best basis for supporting presence. Because each peer knows every other peer on the network then it would be easier for presence information to be distributed to all peers immediately. With such an architecture a central co-ordinator would not be needed. Obviously, as has been discussed elsewhere, this type of architecture suffers in terms of scalability given that a peer would have to broadcast its presence information to every other peer on the network. However if used in small-scale environments then it is the most ideal.

Table 1, suggests possible solutions for satisfying the presence issues with these types of P2P logical architectures.

Issue	Possible solutions for this type of architecture
A	With this architecture you would expect the entity to broadcast its presence information to all peers (and thus all entities also on the abstract level), when it connects to the network.
B	With this architecture you would expect the entity to broadcast the fact that it is disconnecting to all peers (and thus all entities also on the abstract level).
C	In this case it is unlikely that the entity would have informed the rest of the network about its loss of connection. This means that unless peers are expected to regularly broadcast their presence information (and that of any abstract entities they might possess), or periodically poll other peers for theirs, they may be unaware that an entity has been disconnected
D	With this architecture you would expect the entity to broadcast any changes to its presence information to all peers (and thus all entities also

	on the abstract level).
E	Because with this architecture all peers are connected to one another, when a peer (or any abstract entities it might possess) connects to the network, it should automatically discover what other presence information exists on the network.
F	Because with this architecture all peers are connected to one another, when peers (or any abstract entities it might possess) presence information changes, it should automatically inform all other peers (and abstract entities) on the network of this fact. This would not happen, however, if the change were accidental (e.g., loss of power). To take into account this scenario it might be necessary for the individual entities to ping each other at regular intervals.

**Table 1 – Providing presence within direct communication decentralised architectures**

## 5.2 Providing presence within indirect communication decentralised systems

Achieving presence within indirect communication decentralised systems can be difficult (if not impossible) due to the lack of a central co-ordinator. This not only makes it difficult to co-ordinate the collection and publishing of presence information, but also to give peers, users and resources universal id's within the system. Without these id's it is likely to be a difficult to identify what entity's presence information has changed. In addition, because the network can be frequently changing it may be difficult to achieve any form of real-time presence information updates or even for these updates to be received. In theory, when presence information is broadcast onto the network it could be spread between peers using techniques similar to that used for searching or for peer discovery in indirect decentralised architectures. However, due to the dynamic nature of the architecture it cannot be guaranteed that all peers will receive this information, or that a peer that has received it once in the past, will receive it again (due to issues such as network partitioning, alternative message routing, etc).

As a result, it is difficult to analyse the effects these types of architecture can have on the presence issues that have been identified. Table 2, suggests possible solutions for satisfying the presence issues with these types of P2P logical architectures.

<b>Issue</b>	<b>Possible solutions for this type of architecture</b>
A	With this architecture when an entity connects to the P2P network it will be able to publish its presence information but only to those other peers (and their respective abstract entities) it is aware of. In theory, in turn these peers could then broadcast the information to the peers they are aware of.
B	With this architecture when an entity is going to disconnect from the P2P network it will be able to publish this fact but only to those other peers (and their respective abstract entities) it is aware of. In theory, in turn these peers could then broadcast the information to the peers they are aware of.
C	In this case it is unlikely that the entity would have been able to inform the other entities it is aware of, about its loss of connection. This means that unless peers are expected to regularly broadcast their presence

	information (and that of any abstract entities they might possess), or periodically poll other known peers for theirs, they may be unaware that an entity has been disconnected.
D	With this architecture when an entity's presence information changes, it will be able to publish this fact but only to those other peers (and their respective abstract entities) it is aware of. In theory, in turn these peers could then broadcast the information to the peers they are aware of.
E	With this architecture, when the entity connects to the network it would most likely need to perform some sort of presence information discovery. This could operate in a similar manner to resource searching or the discovery of peers within this type of architecture. The entity could issue a presence request that would get passed to all the entities it is aware of. These would return their own presence information, whilst also forwarding the presence request to the entities they are aware of. The issues with such an approach are that it is not overly reliable, returned presence information may not be up to date, the time it takes to gather this information could vary considerably, and it cannot be guaranteed that all entities on the network would receive the request. It is likely that the most reliable presence information would be obtained from the local entities.
F	With this architecture, it could be difficult to keep an entity's presence information up to date with interested parties. Because it cannot be assumed that the entity would be able to contact these parties directly to inform them of a change in the presence information, it is likely that a propagating broadcast method would have to be used. As highlighted previously, this does not provide any guarantees of reliability, and could potentially result in an interested entity not being informed of a change. An alternative strategy would be to make each interested entity attempt to obtain the current presence information itself. However, this could also add to the problems by resulting in the swamping of the network.

**Table 2 – Providing presence within indirect communication decentralised architectures**

One possible solution for supporting presence within indirect communication architectures would be to create small groups of entities. These groups could then communicate with each other in a direct way, but also be linked to other groups. However, to achieve this it is likely that a degree of management would be needed and ultimately this could result in upsetting the equality of the network (in essence turning it into a semi-centralised system).

### **5.3 Providing presence in semi-centralised systems**

It is likely to be easier to achieve presence within semi-centralised based systems than with some of the more decentralised alternatives (in particular with indirect communication decentralised systems) due to the central foci that exist within the system. These foci can be used to capture and publish the presence information that exists, and because all peers within the system will be in direct contact with them, this information should be reasonably consistent and up to date. Obviously the negative side is that these foci become points of failure. Should they go down then this presence information will be lost from the network.

Table 3, suggests possible solutions for satisfying the presence issues with these types of P2P logical architectures.

Issue	Possible solutions for this type of architecture
A	With this architecture you would expect the entity to inform a server node of its presence information when it connects to the network.
B	With this architecture you would expect the entity to inform a server node of the fact that it is disconnecting from the network.
C	In this case it is unlikely that the entity would have informed a server node about its loss of connection. This suggests that as well as presence information updates coming from the relevant entities, the server node(s) may also need to make regular checks on the presence state of the network.
D	With this architecture you would expect the entity to inform a server node of any changes to its presence information.
E	With this architecture, the server nodes would most likely be used to store and distribute presence information around the network. As a result, when an entity connects to the network to obtain current and relevant presence information, it would likely need to send a list of the entities it is interested in, to one of these server nodes. The server node could then return the current presence information for these entities.
F	<p>With this architecture, because the server nodes would most likely be used to store and distribute presence information around the network, they need to be kept informed of what entities are interested in (i.e. who is interested in who).</p> <p>To achieve this each entity would need to register these interests with a server node. In this way, when a change occurs the relevant interested parties can be kept up to date. The main danger with this approach is that amount of information the server may end up having to store, especially if a poor structure is not adopted.</p> <p>An alternative strategy would be to make each entity poll a server node for the current presence information. This, however, could result in swamping the network and server node.</p>

**Table 3 - Providing presence within semi-centralised architectures**

With this architecture an entity could try and obtain current presence information itself by contacting the relevant entities directly. However, this would rely on using previous information about the entities that may have become out of date (peer address, for example).

## 6.0 Discussion – presence design issues

This section attempts to identify some design issues that may need to be considered if it is intended to incorporate presence within a system. These issues include:

- *Controlling presence information*
- *Presence and privacy*
- *Presence consistency*
- *Presence as a mechanism to provide dependability*

## **6.1 Controlling presence information**

Ultimately the presence information that is made available to the rest of the network is controlled by the entity that provides it. Consequently control mechanisms need to be considered and provided.

One possibility is to provide the entity with a range of information levels, ranging from the entity providing very little information to providing a substantial amount. Such an approach can be frequently seen in many distributed applications, where a user can specify as much presence information as they feel (for example, ICQ). This can be further enhanced by also specifying who can have access to this information. For example, all users have access to some presence information, but only my friends have access to all of it.

Anonymity can also be important, as it can mean that an entity can provide presence information but without actually revealing whom it is originating from. For example, an entity can provide 50 MB hard disc space, but it's IP address remains hidden.

Clearly deciding how the presence information is to be captured, is an important issue when considering the use of presence within a system.

## **6.2 Presence and privacy**

To a certain extent privacy is related to how the presence information is controlled. If the control mechanisms are sufficient then privacy should be less of a problem. However there is also the possibility that the presence information that is published may be misused, for example, obtaining a users email address and then using it to send Spam.

Again this may be limited by controlling what information is being made available (i.e. not displaying email addresses). However, in some situations it might be necessary for such information to be made accessible in order for the P2P system to fully function. To reduce the possibility of information misuse in such circumstances it is likely that suitable protection mechanisms would need to be provided. A possible mechanism could be to require authentication before being able to access an entity's presence information. In this way an entity could be sure that only those who have been granted permission (and are trusted) can have access to the information.

If private information is to be published as presence within a system, then it will become important to consider privacy issues, and to reduce the possibility of the information being misused.

## **6.3 Presence consistency**

One important issue with presence information is insuring that it is up to date and valid. Depending on the nature of the system, inaccurate presence information could result in situations, such as, where incorrect business decisions are made.

In order to ensure a high level of validity any updates to the published information needs to be done in a near instantaneous fashion. In reality this is not going to be

achieved, however the type of logical architecture used can play a significant role. In particular, semi-centralised architectures are likely to be the most suitable for distributing quick presence updates around the network.

Ultimately, though, it would need to be decided whether presence mechanisms can be relied on enough, or indeed if they should be, to provide dependable information.

#### **6.4 Presence as a mechanism to provide dependability**

In the Report on the Dependability Properties of P2P Architectures, it was pointed out that many of the identified dependability properties would typically require some form of monitoring mechanisms, in order for them to be properly resolved. Such properties included availability, fault tolerance and maintainability.

Incorporating presence within a system might provide one mechanism with which this monitoring can be achieved. Presence mechanisms could be used to monitor the availability of nodes within the system, or to identify faults that may occur so that the system can act accordingly.

However, although presence could be utilised in this fashion, ultimately it needs to be decided whether presence mechanisms, themselves, can be considered to be reliable, and if not, whether it is possible to make them reliable enough.

### **7.0 Summary**

This document has examined how presence can be supported within P2P systems. It has provided an initial overview of presence and has attempted to identify its main characteristics. It has also examined what requirements need to be satisfied should it be desired to incorporate presence within a system, and how the different logical architectures can have an effect on these.

Initial analysis has shown that semi-centralised architectures would typically provide the best foundation for supporting presence. The server nodes within a semi-centralised system can be used to capture and distribute presence information around the network, and because all nodes connect to the server nodes there are no problems with regards to providing entities with ID's or not being able to contact all parts of the network.

Although, generally, de-centralised architectures are perhaps less suitable for supporting presence, this is not the case for direct communication de-centralised systems. Because all nodes connect to all other nodes, this architecture is arguably the best for supporting presence. Obviously the downside (as reported elsewhere) is that it is not scalable.

This document has concluded by highlighting some issues that may need to be considered if presence support is desired. These have focused on controlling the presence information, ensuring privacy (if desired) and the importance of keeping presence information up to date. The possibility has also been raised of whether presence could be used as a mechanism to improve a systems dependability.

## References

[Godefroid] Godefroid, P., Model Checking for Programming Languages using VeriSoft. *In ACM Symposium on Principles of Programming Languages*, January 1997, 174-186.

[Herbsleb] Herbsleb, J. D., Grinter, R. E., Architectures, coordination, and distance: Conway's law and beyond. *IEEE Software*, Sept/Oct 1999, 63-70.

[Palen] Palen, L., Social, individual, and technological issues for groupware calendar systems. *In CHI'99*, 1999

[Want] Want, R., Hopper, A., Falcao, V., Gibbons, J., The active badge location system. *ACM Transactions on Information Systems*, 10(1):91-102, 1992.