

# A Model for Worldwide Tracking of Distributed Objects\*

Maarten van Steen, Franz J. Hauck, Andrew S. Tanenbaum  
*Vrije Universiteit, Amsterdam*

## Abstract

*We describe a service for locating distributed objects identified by location-independent object identifiers. An object in our model is physically distributed, with multiple active copies on different machines. Processes must bind to an object in order to invoke its methods. Part of the binding protocol is concerned with contacting the object, which offers one or more contact points. A contact point is associated with an active part of the distributed object, and describes exactly how and where initial communication should take place. An object can change its contact points in the course of time, thus exhibiting migration behavior. Finding an object's contact points is the essence of our location service.*

*Our model is based on a worldwide distributed search tree, capable of handling trillions of distributed objects. The tree adapts dynamically to individual migration patterns. By exploiting an object's relative stability with respect to a region, combined with the use of pointer caches, an object can be contacted through a search path of only length two. We present the architecture of our location service, including its update and lookup mechanism, and discuss its scalability.*

## 1 Introduction

Until recently, few people had any conception of a worldwide distributed system. To them, worldwide communication just meant communication through telephones, faxes, electronic mail, etc. The introduction of the World Wide Web and the ease of access to the Internet is now radically changing our perception of worldwide distributed systems. Such systems should allow us to easily share and exchange information. This also means that it should be easy to track sources of information, even if these sources move between different locations. (We do not address the problem of finding *relevant* sources of information as is done by, for example, resource discovery services [13].) A key role in tracking sources of information is played by naming systems.

An important problem with current naming systems for wide area networks is that names are location-dependent: a name is tightly coupled to the location of the object it refers to. For example, a URL such as `http://www.ripn.net/nic/rfc/rfc1737.txt` is the name of a Web page containing the text of RFC 1737. The name reflects exactly where the page is stored. If the page is moved or replicated, the name will have to change as well. In general, we have the following:

- Location-dependent names make it difficult to handle migration. Migrating an object cannot be done without either changing its name or making the name become a forwarding reference. In wide area systems, the latter may lead to long chains of references, which are inefficient and susceptible to node and link failures.

---

\*A shorter version of this paper appeared in the Proceedings of TINA'96, Heidelberg, September 1996.

- Location-dependent names cannot be used easily in combination with replication. Clients must keep track of which object is replicated, where the replicas are located, and how consistency between these replicas is maintained. These issues, however, should be hidden entirely from clients.

What is needed is a naming and identification facility that hides all aspects of an object's location. Users should not be concerned where an object is located, whether it can move, whether it is replicated, and if it is replicated, how consistency between replicas is maintained. This mechanism should be available to all applications as a standard facility. Above all, it should scale to the entire world, and be able to handle trillions of objects.

In this paper, we present a model for locating objects using location-independent identifiers. We use a search tree in which an object's contact addresses are stored at relatively stable locations, near to the places where the object can be reached. We show how these stable locations are identified dynamically, and that they may change as the migration behavior of the object changes over time. Storing contact addresses at stable locations permits us to effectively cache location pointers. The combination of dynamically identifying stable storage locations for contact addresses, and caching pointers to those locations, is new. The result is a location service that is highly efficient by exploiting locality in lookup and update operations.

The paper is organized as follows. In Section 2, we present an architecture of an object-based distributed system that is used as a framework for our discussion of locating objects. In Section 3, we explain the logical organization of a scalable location service, followed, in Section 5, by a presentation on how such a service can be implemented. We conclude with a discussion and comparison to related work in Personal Communication Networks and distributed object systems in Section 6.


## 2 System architecture

In this section we present the overall system architecture for object support. We introduce the concept of a distributed shared object and briefly explain the role of naming and location services in our system.

### 2.1 Distributed shared objects

In our model, processes interact and communicate through **distributed shared objects** [6, 18]. Each object offers one or more interfaces, each consisting of a set of methods. Objects are passive, but multiple processes may access the same object simultaneously, allowing them to communicate by reading and changing the object's state through method invocations. Changes to the object's state made by one process are visible to the others. A major distinction with other object-based models is that objects are physically distributed, meaning that active copies of an object's state can and do reside on multiple machines at the same time. However, processes are not aware of this: state and operations on that state are completely encapsulated by the object. This means that all implementation aspects, including communication protocols, replication strategies, and distribution and migration of state, are part of the object but are hidden behind its interface. This model is illustrated in Figure 1.

In order for a process to invoke an object's method, it must first bind to that object. This means that an interface belonging to the object, as well as an implementation of that interface is to be placed in the process' address space. Such an implementation is called a **local object**. A local object is somewhat comparable to an RPC stub [2] or a proxy [15]. However, local objects are more powerful: they allow for different implementations of the same interface. For example, one process may just want a proxy



**Figure 1:** Example of an object with one contact point and distributed over three address spaces.

implementation, whereas another may want an active copy of the distributed object’s state. Such differences are transparent to the using processes: they only see the interface to the distributed object. This is an important distinction with most other object-based models, which generally follow a client/server approach such as DCE [14], CORBA [12], and Spring [10]. In those models, an object’s state is located entirely in a single address space and is remotely accessed through proxies. These models do not support replicated state as we do. Instead, replication is applied to an entire object. A separate systemwide replication policy is used to keep the replicas consistent, and to handle communication with clients.

As an example of our approach, consider the USENET news system. Suppose we make each newsgroup an object. When a process is interested in a particular newsgroup, it binds to the corresponding newsgroup object, loading a local newsgroup object into its address space. The process can then get and post articles without having to worry about which articles are kept where. The newsgroup object may support any of several replication policies, ranging from having all news at a central repository to having every article stored on every machine. The object independently determines its replication policies, and encapsulates all implementation aspects.

Participating in an object’s implementation requires that a process first gets in touch with the object. To this end, a distributed object has one or more contact points. A **contact point** specifies the network address and protocol by which initial communication with the object can take place. A contact point is bound to a local object of the distributed object, as also shown in Figure 1. For example, a distributed object may offer a contact point at network address *netaddr* to be used for a TCP/IP connection. It may be further specified that a process that wants to initially contact the object should use an object-specific protocol, say *initprot*. A **contact address** describes a contact point and consists of two components: (1) the protocol specification (“use *initprot* over TCP/IP”), and (2) the network address (“*netaddr*”). A client will then first have to install a local object that implements *initprot*, and subsequently tell it to use *netaddr* to contact the object.

An object’s contact points may change in the course of time. For example, an object can be said to expand to, or withdraw from a region when contact points in that region are established or removed, respectively. Finding an object’s contact points is the first step of binding, and is the main subject of

this paper.

## 2.2 Naming and location services

In order to find contact points, we propose a two-level naming hierarchy. The first level deals with hierarchically organized, user-defined name spaces. These name spaces are handled by a distributed **naming service**. However, where traditional name service implementations maintain name-to-address bindings, names in our approach are mapped to object handles. An **object handle** is a globally unique, and location-independent object identifier. Object handles form the second level in the naming hierarchy. Each object handle is mapped to an object's contact addresses. It is the purpose of a **location service** to maintain the mapping between object handles and contact addresses.

For our system, we have identified that the location service should meet the following main requirements.

**Scalability.** The service should allow clients and objects to be located anywhere in the world, and be able to support a huge number of objects. We anticipate that eventually, one billion users with 1000 objects each will be registered with the location service, adding up to  $10^{12}$  objects.

**Locality.** Assuming that the cost for looking up a contact address generally decreases with the length of the route to the corresponding contact point, we require that the location service exploits locality. This means that if one of an object's contact points is near to the place where a lookup request is issued, it should be easy to find the object. Likewise, adding or deleting a contact point should be relatively cheap in a region where an object already has many contact points.

**Stability.** Objects differ with respect to their migration patterns, expressed as updates on their set of contact addresses. An object is said to be **stable** with respect to a region, if a subset of its contact points generally lie in that region. If an object is stable with respect to a region  $R$ , we require that searching or updating the object's contact addresses in  $R$  is cheaper than when the object is not stable with respect to  $R$ .

**Fault tolerance.** The location service should be resilient to node and link failures, and at the least, should degrade gracefully in terms of performance and functionality.

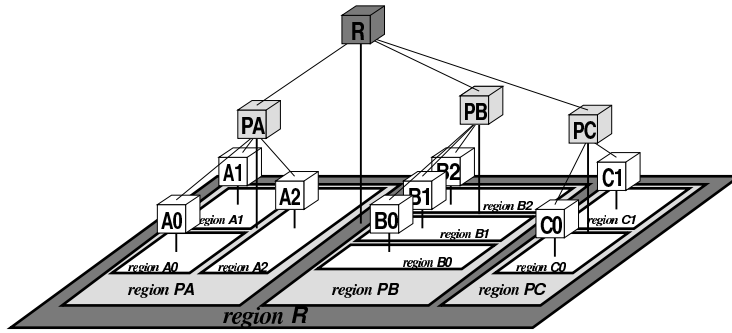
Location services are not new and have shown to be relatively easy to implement in local distributed systems. However, they become much more complicated when scalability is taken into account as we discuss next.

## 3 The location service

In our model, we assume a hierarchical decomposition of a (worldwide) distributed system into regions. For example, a lowest-level region may consist of a departmental local area network, whereas a region one level higher may constitute a campus-wide network. This decomposition is relevant to only the location service. It is entirely transparent to client processes.

With each region we associate a **directory node**, capable of storing contact points that lie within that region. This leads to a logical organization as shown in Figure 2. As we explain in Section 5, this organization corresponds to a *virtual* search tree, in the sense that each directory node is distributed across several physical nodes for scalability purposes.

Each lowest-level region contains one or more distributed objects called **location resolvers**. A location resolver offers an interface to the location service, making the latter appear as just another

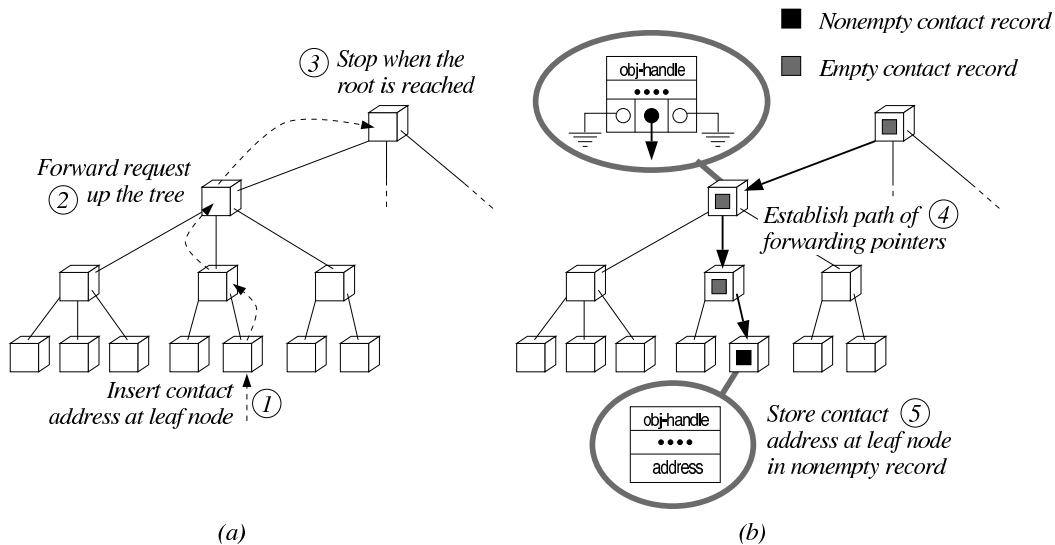


**Figure 2:** The logical organization of the location service as a virtual search tree.

distributed shared object. Additionally, a location resolver encapsulates locality by allowing clients to transparently communicate with the leaf node for its region. A process is automatically bound to a location resolver.

### Inserting and deleting contact addresses

A contact address is initially entered and stored at the leaf node of the tree representing the location of the corresponding contact point. The location service also maintains, per object, a path of forwarding pointers from the root to each leaf node where a contact address is stored. Contact addresses and forwarding pointers are stored in **contact records**. An empty contact record contains only forwarding pointers, whereas a nonempty contact record will contain at least one contact address. An implication of this design is that in the worst case, it is always possible to locate every object by following the chain of pointers from the root node. In practice, we can do much better than this, as described later.

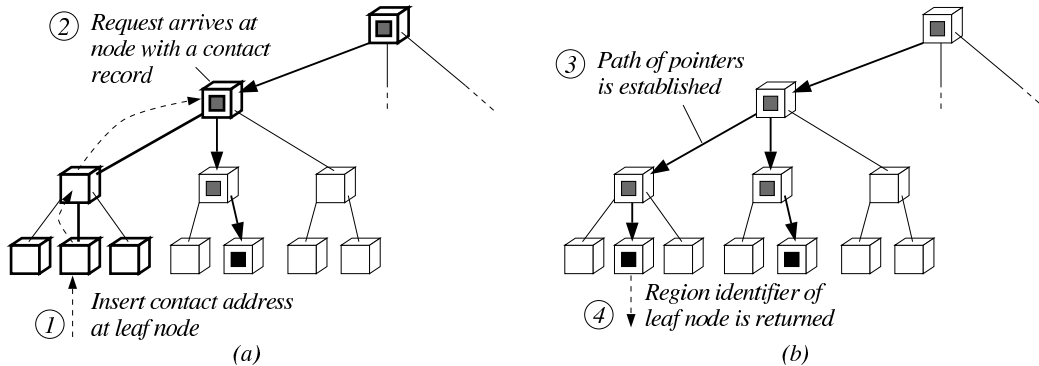


**Figure 3:** Inserting a contact address for a previously unregistered object, resulting in a path of forwarding pointers from the root to the leaf node.

Figure 3 illustrates the insertion of the first contact address for an object. First, the request for insertion is propagated up the tree to the root, as shown in Figure 3(a). Then, a path of forwarding

pointers is established from the root to the leaf node where the insertion takes place. An empty contact record containing a forwarding pointer is created at each intermediate node, as shown in Figure 3(b). The address is finally stored in a record in the leaf node.

When a part of the path already exists, for example when inserting a second address in a different region, only the missing pointers are established. As shown in Figure 4(a), an insertion request propagates upwards to the first higher-level directory node where the object is already known. From there on, a path of forwarding pointers is established to the leaf node where the insertion takes place, as shown in Figure 4(b). In the case that there is already a contact record for the object at the leaf node, the new address is simply added to that record.



**Figure 4:** Inserting a contact address when the object is already known. Only the missing pointers are established.

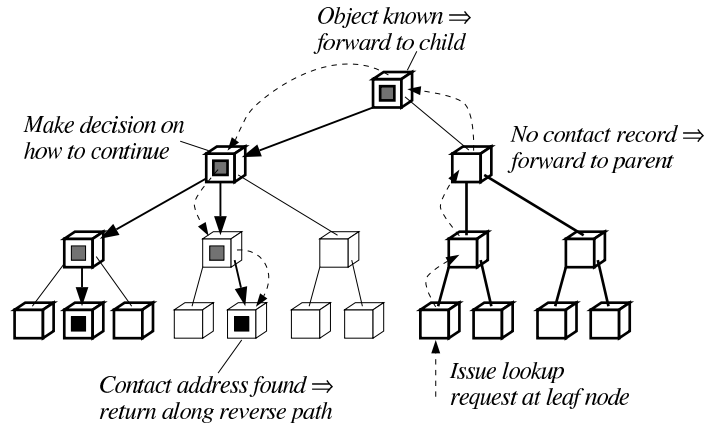
The insert operation returns a **region identifier** identifying the leaf node where the insertion takes place, and which can be used for deletion. Deleting a contact address is straightforward and is done as follows. First, the address is found through a search path up the tree, starting at the leaf node identified by the region identifier that was returned when the address was inserted. We also support deletion of contact addresses without knowing the region identifier, by an exhaustive search through the tree. Once the contact address has been found, it is removed from its record. If a contact record contains no contact addresses or forwarding pointers, it is deleted. The parent directory node is informed that it should delete its forwarding pointer to that record, possibly leading to the (recursive) deletion of the object's contact record at the parent node.

It is seen that inserting and deleting contact addresses exploits locality, especially when contact addresses already exist in the region where the operation is being performed.

### Looking up contact addresses

Looking up a contact address is done as follows. A process provides its location resolver with an object handle, a quasi-random binary number, which it will often have obtained by means of the naming service. The handle is then given to the leaf node of the resolver's region. As shown in Figure 5, a linear search path is established starting at the client's leaf node, and upwards to the first directory node where the object is already known. In the worst case, this means propagating the request up to the root. The path then continues downwards to a leaf node, whose contact addresses are then returned to the requester.

Note that there may be alternative ways for continuing a search path when a directory node has several forwarding pointers to different contact records. Strategies for selecting amongst alternatives



**Figure 5:** The default approach for looking up a contact address.

are beyond the scope of this paper. Furthermore, clients can specify a minimum and maximum number of contact addresses that should be returned. The location server then applies backtracking to find contact addresses. When a nonempty contact record is reached, its contact addresses are aggregated, but the search continues if the minimum number has not yet been reached.

Again, it is seen that locality is exploited: the lookup operation searches local regions first, and gradually expands to larger regions if no contact addresses are found.

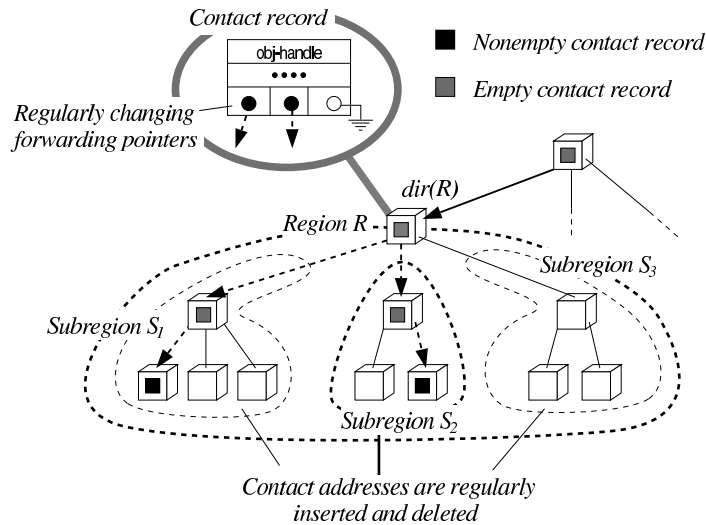
## 4 Optimizations

We now explain how the location service dynamically finds optimal solutions on a per-object basis. The optimizations are based on the observation that caching mechanisms are only truly effective if cache entries change infrequently. For the location service, the only data it has full control over are the placement of contact addresses in contact records (not the placement of the objects themselves). This means that if we can place contact addresses in stable locations, we can make effective use of pointer caches during lookup operations. As we show below, we can even do this for highly mobile objects. Stabilizing the placement of contact addresses and subsequently constructing pointer caches, is a distinctive and novel feature of our approach.

### Stabilizing the placement of contact addresses


Our first concern is to decide in which contact record a contact address is to be stored. By default, an object's contact address is stored in its contact record at the leaf node where it was initially inserted. Now consider some region  $R$  as shown in Figure 6, and assume that an object  $O$  is changing its contact addresses regularly between the subregions  $S_1$ ,  $S_2$ , and  $S_3$ . Also, assume that there is always at least one contact address somewhere in  $R$ , so that there will always be a contact record for  $O$  at directory node  $dir(R)$ .

Each time the object expands to a subregion  $S_k$  the location service creates a path of forwarding pointers from  $dir(R)$  to a leaf node in  $S_k$ . Likewise, when withdrawing from  $S_k$  the path has to be deleted. If expansion and withdrawal occurs regularly, it makes sense to store the contact address in the object's contact record at  $dir(R)$ , thus saving the cost of path maintenance. In addition, addresses



**Figure 6:** The situation of an object regularly moving between subregions.

of contact points in any of the subregions are now stored in a stable place, namely at the directory node  $dir(R)$ . This leads to a situation shown in Figure 7.



**Figure 7:** Storing contact addresses in a stable place at a higher level.


Of course, the migration behavior of an object with respect to a region may change. For example, assume the contact record at  $dir(R)$  has contained a contact address for subregion  $S_k$  for quite some time. In that case, the contact address will be propagated to a directory node in  $S_k$ , because apparently, stability occurs in a smaller region than  $R$ .

Stability is measured by timestamping contact addresses and forwarding pointers, as well as recording how long an object has not had a contact point in a region. In all cases, history is taken into account by weighted accumulation of old and new timing information.

### Using pointer caches

By storing contact addresses in stable contact records, our model leads to the construction of a search tree *per object*, in which contact records tend to remain in place even if the associated object moves.

This permits us to effectively shorten search paths by caching *pointers* to contact records. Specifically, a pointer to the directory node containing a nonempty contact record is cached at each node of the search path when returning the answer to the leaf node where a lookup request originated as shown in Figure 8.



**Figure 8:** Installing cache pointers after looking up a contact address in the tree from Figure 7.

The combined effect of pointer caches and stable contact records should not be underestimated. An object that primarily moves within a region  $R$  can be tracked by just two successive lookup operations: the first one at the leaf node servicing the requesting process, and the second one at the directory node for region  $R$ . Moreover, our solution forwards a request in the direction of a contact point. This is a considerable improvement over existing approaches.

Caches are kept consistent in a lazy fashion. A timeout value is associated with each cache entry, which depends on the stability of the referenced contact record. Caches can then, for example, be regularly purged by a sweep algorithm, thus preventing the use of timers per cache entry. In principle, the only other time that a cache entry can be invalidated is when the referenced contact record has been deleted, or when it no longer contains contact addresses.

In the first case, the cache entry is removed and a lookup follows the default strategy described before. In the second case, several strategies are possible. One is to continue the lookup from the directory node containing the referenced contact record. Once contact addresses have been found, the search path is traversed in the opposite direction, thereby updating or establishing cache entries of intermediate nodes. Alternatively, the cache entry can be immediately invalidated after which the default lookup strategy is followed.


## 5 Scalability and implementation

The search tree described so far obviously does not scale. In particular, higher-level directory nodes not only have to handle a relatively large number of requests, they also have enormous storage demands. For example, the root node needs to maintain a contact record for *every* registered object. This requires already a storage capacity in the terabyte range. In this section we show that our location service scales using a partitioned implementation of the search tree.

### Partitioning directory nodes

We partition a directory node into one or more **directory subnodes**, such that each subnode is responsible for a subset of objects. As an example, we can use the first  $n$  bits of an object handle to identify the subnode responsible for that object. Subnodes of a particular directory node need not communicate with each other since they maintain different subsets of objects, and all operations are performed on a

per-object basis. Communication between directory nodes in the original search tree only takes place between their respective subnodes. To illustrate, Figure 9 shows a search tree in which the root node has been partitioned into four subnodes based on the first two bits ( $n = 2$ ), and each of the leaf nodes into two subnodes ( $n = 1$ ).



**Figure 9:** A search tree and a corresponding logical tree after partitioning the directory nodes into subnodes.

In this example, we can guarantee a reasonable load balance if the object handles are uniformly distributed, or otherwise, if the partitioning is based on uniformly distributed numbers, uniquely derived from object handles. This is the case, for example, if the first  $m$  bits of an object handle are always generated as a random number and  $n \leq m$  for all  $n$ .

In principle, each directory node can be partitioned independently according to the number of available hosts and the expected load. In practice, this independence can compromise scalability: each parent subnode may need to maintain a link to every child subnode. Consequently, the number of links between a partitioned parent node and its partitioned children may be prohibitively large. In our example, this problem does not occur. By agreeing to base *all* partitions on the *leftmost* bits of object handles, we have chosen for a common partitioning scheme. However, each node may still individually decide on the number of leftmost bits it will use for partitioning. Without going into further details here, it can be shown that such partitioning schemes can be readily devised. The result is that the number of parent-child links is dramatically reduced and that scalability is not compromised.

Finally, if we assume that the root can be partitioned into 10,000 subnodes (which means that a single root subnode will have to provide service to approximately  $10^9/10^4 = 100,000$  users), we need also not run into any storage problems. If an object has, on average, five contact points, a contact record at the root will consist of five pointers, the object's handle, and some additional stability information. We expect that this requires no more than 100 bytes. With a total of  $10^{12}$  objects, each of the 10,000 root subnodes will have to store 10 gigabytes of data. This is definitely manageable. Similar results can be derived for the other directory nodes.

### Implementing directory subnodes

As communication between directory nodes in the original search tree now takes place between their respective subnodes each subnode should be aware of how the directory node with which it communicates is actually partitioned. This information is contained in what is called a **meta node** of which there is one per directory node. A meta node also maintains the mapping of subnodes to physical nodes.

Each subnode of a directory node knows where the respective meta nodes of its parent and children can be reached. We assume that the mapping of meta nodes onto physical nodes is reasonably stable. Partitioning and mapping information contained in a meta node is also assumed to be relatively stable,

so that it can be easily cached by subnodes. This assumption is necessary to avoid that meta nodes are queried each time a subnode needs to communicate with its parent or children, which would turn the meta node into a potential communication bottleneck.

## 6 Discussion and related work

We have made a strict separation between a naming service which is used to organize objects in a way that is meaningful to their users, and a location service which is strictly used to contact an object given a unique identifier. Naming services can be used for finding information based on the *meaning* of a name, as is often used for Internet resource discovery services [13]. In our scheme, information retrieval would start with finding relevant names, retrieving the associated object handles, and having the location service return contact address for each object that was found to be potentially interesting.

Location services are particularly important when sources of information, i.e. objects, can migrate between different physical locations. They are becoming increasingly important as mobile telecommunication and computing facilities become more widespread. To relate our work to that of others, we therefore concentrate primarily on aspects of mobility, for which we make a distinction between mobile hosts and mobile objects.

### Mobile hosts

So far, much research has concentrated on *mobile hosts*, usually either hand-held telephones or mobile computers. A characteristic feature of these hosts is that their mobility is directly coupled to that of their user. This has two important consequences that do not apply to mobility as considered in this paper. First, the speed of migration is limited to the maximum speed at which a person can move: about 1000 miles per hour. Second, a host is always at precisely one location. There is no notion of multiple contact points as we have introduced in our model.

Both features influence the design of a location service. By assuming a speed limit with respect to migration, it becomes possible to adopt a strategy in which data structures gradually adapt as the object moves. This has been used in very different types of location services. For example, mobile hosts in the Internet are supported at the network level where routing tables are dynamically adapted as a host moves [11]. In location services which make use of a distributed search tree, higher-level nodes can effectively cache network addresses rather than pointers to addresses. Whenever a host moves, it leaves a forwarding pointer to its next location. Caches are updated when a considerable distance has been traveled (as described in [1]), or after a lookup. But the speed limitation may even make it possible to avoid the use of caches altogether as illustrated in [19].

### Mobile objects

The situation becomes entirely different when dealing with mobile objects that (1) are not statically bound to a single host, and (2) can travel at almost the speed of light. This is the case, for example, with World Wide Web pages and distributed (shared) objects in general.

Mobile objects have mainly been considered in the context of local distributed systems. In Emerald [8], mobile objects are tracked through chains of forwarding pointers, combined with techniques for shortening long chains, and a broadcast facility when all else fails. Such an approach does not scale to worldwide networks. An alternative approach to handle worldwide distributed systems is the Location Independent Invocation (LII) described in [3]. By combining chains of forwarding references, stable storages, and a global naming service, an efficient mechanism is derived for tracking down objects.

Most of the applied techniques are orthogonal to our approach, and can easily be added to improve efficiency. However, the global naming service which is essential to LII assumes that update-to-lookup ratio is small. Designing a global location service that is not based on such an assumption is an important goal of our research.

A seemingly promising approach that has been advocated for large-scale systems are SSP chains [16]. The principle has been applied to a system called Shadows [4]. SSP chains allow object references to be transparently handed over between processes. In essence, an object reference is just a pointer to a data structure (called a *stub*), that acts as a representative for the object. When passing an object reference from one process  $P$  to another process  $Q$ , a network connection between the two is established.  $P$ 's endpoint of this connection is called a *scion*, that of  $Q$  is the stub just mentioned. The scion is capable of resolving an incoming reference to the referenced object, possibly by passing it to the stub in its own address space. Consequently, there is no need for any location service because an object reference can always be resolved through the chain of scion-stub pairs that the holder of the reference has established with the object. A serious drawback of this approach is that exploiting locality is completely neglected. This is completely unacceptable for worldwide systems.

### Contributions of our approach

One of the main advantages of our approach is that our location service can handle distributed objects that have several contact points and that show arbitrary migration patterns. In contrast to the approach described in [9], we do not adapt update and search strategies to migration patterns, but adapt the search tree on a per-object basis instead. By registering contact points in the smallest region in which (part of) the object is moving we can make effective use of pointer caches. The combined effect is an extremely short search path, in the optimal case of only length two, from a client to the object. In just two hops it is possible to locate even seemingly randomly migrating objects. This is a considerable improvement over existing approaches.

The use of pointer caches instead of data caches has also been proposed for PCNs. The main reason to apply caching in those cases is to avoid excessive network traffic to the *home location* of a host, which forms the root of a two-level search tree. Caching is done at the second level, by pointing to locations where the host is expected to be found. Cache consistency is achieved by invalidation on demand [7], but can even be done through active updates [17]. However, caches in PCNs do not account for update patterns. As also observed in [5], exploiting locality in location updates can reduce tracking costs. A distinctive feature of our approach compared to PCNs, is that we have several levels allowing us to exploit locality more effectively by inspecting succeedingly expanded regions at linearly incrementing costs. On the other hand, locality is also exploited in location updates, making our pointer caches highly effective.

### References

- [1] B. Awerbuch and D. Peleg. "Online Tracking of Mobile Users." *J. ACM*, 42(5):1021–1058, 1995.
- [2] A.D. Birrell and B.J. Nelson. "Implementing Remote Procedure Calls." *ACM Trans. Comp. Syst.*, 2(1):39–59, 1984.
- [3] A. Black and Y. Artsy. "Implementing Location Independent Invocation." *IEEE Trans. Par. Distr. Syst.*, 1(1):107–119, 1990.
- [4] S. Caughey and S.K. Shrivastava. "Architectural Support for Mobile Objects in Large-Scale Distributed Systems." In L.-F. Cabrera and M. Theimer, (eds.), *Proc. Fourth Int'l Workshop on Object Orientation in Operating Systems*, pp. 38–47, Lund, Sweden, Aug. 1995. IEEE.
- [5] J.S.M. Ho and I.F. Akyildiz. "Local Anchor Scheme for Reducing Location Tracking Costs in PCNs." In *Proc. First Int'l Conf. on Mobile Computing and Networking*, Berkeley, CA., Nov. 1995. ACM.

- [6] P. Homburg, L. van Doorn, M. van Steen, A. Tanenbaum, and W. de Jonge. "An Object Model for Flexible Distributed Systems." In *Proc. First ASCI Annual Conf.*, pp. 69–78, Heijen, The Netherlands, May 1995.
- [7] R. Jain, Y.-B. Lin, and C. Lo. "A Caching Strategy to Reduce Network Impacts of PCS." *IEEE J. Selected Areas Commun.*, 12(8):1434–1444, 1994.
- [8] E. Jul, H. Levy, N. Hutchinson, and A. Black. "Fine-Grained Mobility in the Emerald System." *ACM Trans. Comp. Syst.*, 6(1):109–133, 1988.
- [9] P. Krishna, N.H. Vaidya, and D.K. Pradhan. "Location Management in Distributed Mobile Environments." In *Proc. Parallel and Distributed Information Systems*, pp. 81–88. IEEE, 1994.
- [10] J. Mitchell et al. "An Overview of the Spring System." In *Proc. Comcon Spring 1994*. IEEE, Feb. 1994.
- [11] A. Myles and D. Skellern. "Comparing four IP based Mobile Host Protocols." *Computer Networks and ISDN Systems*, 26(3):349–355, 1993.
- [12] Object Management Group. "The Common Object Request Broker: Architecture and Specification, version 1.2." Technical Report 93.12.43, OMG, 1993.
- [13] K. Obraczka, P.B. Danzig, and S.-H. Li. "Internet Resource Discovery Services." *Computer*, 26(9):8–22, 1993.
- [14] W. Rosenberry, D. Kenney, and G. Fisher. *Understanding DCE*. O'Reilly, Sebastopol, Calif., 1992.
- [15] M. Shapiro. "Structure and Encapsulation in Distributed Systems: The Proxy Principle." In *Proc. Sixth Int'l Conf. on Distributed Computing Systems*, Boston, MA, May 1986. IEEE.
- [16] M. Shapiro, P. Dickman, and D. Plainfossé. "SSP Chains: Robust, Distributed References Supporting Acyclic Garbage Collection." Technical Report 1799, INRIA, Rocquencourt, France, 1992.
- [17] N. Shivakumar and J. Widom. "User Profile Replication for Faster Location Lookup in Mobile Environments." In *Proc. First Int'l Conf. on Mobile Computing and Networking*, Berkeley, CA., Nov. 1995. ACM.
- [18] M. van Steen, P. Homburg, L. van Doorn, A.S. Tanenbaum, and W. de Jonge. "Towards Object-based Wide Area Distributed Systems." In L.-F. Cabrera and M. Theimer, (eds.), *Proc. Fourth Int'l Workshop on Object Orientation in Operating Systems*, pp. 224–227, Lund, Sweden, Aug. 1995. IEEE.
- [19] J.Z. Wang. "A Fully Distributed Location Registration Strategy for Universal Personal Communication Systems." *IEEE J. Selected Areas Commun.*, 11(6):850–860, 1993.