Spheres: A New/Old Model for Distributed Computing

Mark Miller and Christian Scheideler
Dept. of Computer Science
Johns Hopkins University
3400 N. Charles Street
Baltimore, MD 21218, USA
{markm,scheideler}@cs.jhu.edu

November 19, 2004

Abstract

Peer-to-peer systems have recently attracted an enormous amount of attention both inside and outside of the research community. Systems such as Gnutella and KaZaA have demonstrated that distributed applications are feasible even when spanning millions of potentially unreliable sites, which has led to a renaissance in distributed systems research in recent years. Interestingly, there seems to be a common theme underlying all of the recent peer-to-peer system designs which significantly differs from classical distributed computing approaches, but it has apparently not been formalized in the distributed computing community so far, though the basic ideas are so old that they predate the distributed computing era. We attempt to formalize this theme by proposing a so-called *Spheres* model that, as we will argue, is simple to use and yet provides a universal platform for the development of efficient and robust distributed systems including, of course, peer-to-peer systems.

1 Introduction

Writing distributed programs is a challenging task. Many researchers have tried to overcome this problem with the help of a suitable programming environment, using message passing in the 1970s [1, 14, 25], remote procedure calls in the 1980s [11, 24, 35], and distributed shared memory (DSM) in the 1990s [2, 3, 22, 29]. In fact, there has been a vast amount of projects on platforms for distributed computing. See `http://www.computer.org/parascope/` or `http://www.aspenleaf.com/distributed/` for a comprehensive list. Still, the percentage of distributed applications based on these environments is very small. One possible explanation for this is that each round is an evolutionary stage for the distributed programming paradigm, and each attempt brings us closer to a methodology that will ultimately be easy to use and efficient. A less optimistic explanation of the failure of each attempt is that no matter what methodology people will come up with, ease of use and efficiency are two concepts that cannot be achieved at the same time [34]. We believe that this is in fact true for the current direction in which the design of universal platforms for distributed programs is currently heading. The basic problem is that all of these platforms are more or less following the von Neumann approach of separating code from data, but this is the *wrong* approach for providing a universal paradigm for distributed computing. But if the von Neumann approach is the wrong approach, is there any right approach?

We believe that the right direction is a paradigm following the peer-to-peer approach. Peer-to-peer systems have recently attracted an enormous amount of attention both inside and outside of the research community. Systems such as Gnutella and KaZaA have demonstrated that distributed applications are feasible even when spanning millions of potentially unreliable sites, which has led to a renaissance in distributed systems research in recent years. Many peer-to-peer systems have since been implemented in the academic community. Among them are Chord [33], CAN [31], Pastry [17], Tapestry [36], GISP [21], CIRCLE [13], Bamboo [9], and MLDonkey [27], to name just a few, as well as more open architectures such as I3 [32], OpenHash [26], KBR [16], and PeerWare [15]. Further links to middleware projects for peer-to-peer, distributed, grid, and collaborative computing can be found at `http://carstad.gsfc.nasa.gov/Topics/JBRESEARCH/PeerToPeer.htm`. Interestingly, there is a common theme underlying these peer-to-peer systems that is much more flexible and powerful than the von Neumann approach, but it has apparently not been formalized in the distributed computing community, though the basic ideas behind it are so old that they predate the distributed computing era. We attempt to formalize this theme by proposing the so-called *Spheres* model that, as we will argue, is simple to use and yet provides a universal platform for the development of efficient and robust distributed systems. But before we present this model in detail, we start with some basic thoughts about a paradigm useful for distributed computing.

1.1 Towards a useful paradigm for distributed computing

Any paradigm that claims to be useful for distributed computing must be acceptable by all groups involved: users, developers, and scientists. This means that it has to satisfy three central demands:

- It must be *easy* to apply,
- it must allow the development of *efficient* distributed programs, and
- it must be *simple and precise* to allow a verification and formal analysis of these programs.

Though in the academic world, ease of use may not be the most important issue, it should be clear that no matter how good a paradigm is, if it requires an expert to apply, it will not gain wide-spread acceptance. Also, a paradigm that does not allow the development of efficient distributed programs will most likely not be used for anything else than prototyping, and will therefore not make the transition from academia to industrial applications.

On the other hand, any programming paradigm that claims to allow the development of efficient distributed programs must take the following issues into account:

- Sites operate in an asynchronous environment,
- sites may join and leave the system, or may simply fail,
- sites have different resources (processing cycles, memory, bandwidth), and
- messages have varying delays, or may simply get lost.

Thus, distributed applications should be given a high degree of freedom to manage their resources, which seems to forbid a paradigm which is easy to apply and precise. On the other hand, the freedom given to the developer should not be so high that it is tempting to produce inefficient code rather than efficient code. Thus, besides the paradox of achieving ease of use and preciseness together with a high degree of freedom at the same time, we also have to fight with the paradox of offering a high degree of freedom and restricting the development of inefficient code at the same time. Can there possibly be a paradigm that resolves these paradoxes?

If so, we are searching for a paradigm that needs as few constructs as possible while offering a maximum amount of flexibility and avoiding major pitfalls in the development of inefficient code. To fulfill the latter two of these demands, a useful paradigm should

1. allow a distributed application to control the placement of processes and data as well as their degree of reliability,
2. introduce a hidden overhead that is as low as possible, and
3. enforce that processes in a distributed application are decoupled in space, time, and flow.

Space decoupling means that the interacting processes do not need to know their physical locations, *time decoupling* means that the interacting processes do not need to be actively participating in the interaction at the same time, and *flow decoupling* means that the code execution inside processes is not blocked by outside interactions.

Standard distributed computing approaches such as message passing, remote procedure calls, and shared spaces can only provide decoupling for a subset of these issues [18]. So a different approach is needed.

1.2 Why we should not view the network as a von Neumann machine

We all know how to write programs for a single computer. In doing so, we are usually following the von Neumann paradigm without being explicitly aware of it: Code and data are separate entities. The data is stored in a linear addressable and reliable memory and can be accessed at unit cost. So wouldn't it be a good idea to provide a platform for distributed applications that emulates the properties of a von Neumann machine, i.e. that provides a shared space and allows to establish processes acting on this shared space? The problem with this approach is that a universally applicable distributed shared memory cannot be implemented efficiently.

A major cause of this inefficiency is the communication overhead associated with managing the shared space in a scalable and reliable fashion. But even if these problems can be overcome by a shared memory environment, the mere attempt of hiding networking issues may create the illusion to the programmer that there is a single, natural design for a given application, regardless of the context in which that application will be deployed. This assumption can lead to the design of parallel programs that perform poorly in a distributed environment by using, for example, a parallelism that is too fine-grained to hide the network latency, or a communication pattern that does not reflect the underlying communication network. Other critical issues are:

1. Processes and data are *separate* entities. This creates access and sharing problems that have to be handled with great care to avoid inefficiencies and inconsistencies.

2. Read and write accesses can potentially *block*.

3. There is a *linear addressable* memory. This makes it tempting to use implicit and/or non-local strategies such as parallel hashing [20, 10] that are (or appear to be) efficient in theory (when using the unit-cost access assumption) but that are actually expensive to emulate in a distributed environment.

4. The memory is supposed to be *reliable*. This has made it tempting in the past to design distributed data structures in which only the *accesses* are parallelized but not the *structure* itself, creating data structures that are vulnerable in a distributed setting [19, 30].

5. The memory management (mapping and access control) is *hidden* from the application. This prevents an application from optimizing the data placement to its needs.

6. It is not specified how processes can join an application. So this has to be handled by hand in distributed applications, making their design more complicated.

Finally, the shared memory concept can create security problems. In server-based systems, it may be possible for parts of a distributed application to run on top of a virtual address space without sacrificing memory protection: different parts of the application execute in private protection domains controlled by the server, and segments are shared only by mutual agreement. However, in open, server-less systems such as peer-to-peer systems, it may be impossible for the application platform to decide which part of a distributed application should have the exclusive right of using a particular memory segment, as long as the shared memory is a priori open to anyone.

The deficits above reveal that we have to make major changes if we want to arrive at a paradigm useful for distributed applications:

1. Instead of separating processes and data, organize all computation and storage into so-called *spheres*, autonomous, atomic threads with their own, private processing and data.

2. Instead of read and write operations, that may potentially block, only allow strictly non-blocking operations.

3. Instead of a linear addressable memory, there is no memory besides the private memory within spheres, and spheres are interconnected by an explicit link structure which is under the control of the application.

4. Instead of assuming reliable resources, spheres may fail.

5. Instead of a hidden resource management, the sphere placement and migration is under the control of the application.

6. Instead of a manual integration of new spheres, provide a rendezvous mechanism that allows spheres to find each other in an automatic way.

In the next section we present a formal framework that takes these insights into consideration, and we will argue why we believe that this framework is useful for theory and practice.

2 The Spheres Framework

In this section, we suggest a formal framework for the development of distributed applications. The basic ideas behind this framework date back to the actors model developed by Carl Hewitt at MIT in the area of artificial intelligence [12], at a time when distributed computing was still in its infancy. Unfortunately, software and hardware issues at that time prevented his ideas from becoming wide-spread. So people rather followed the von Neumann paradigm and its variants (including models such as BSP, logP, QSM, HMM, and many others).

2.1 The spheres framework

The spheres framework consists of three layers:

- **Network layer**: this is the lowest layer. It handles the exchange of messages between peers.
- **Contact layer**: this handles the exchange of messages between spheres and allows spheres to locate and interconnect to each other.
- **Spheres layer**: this is the layer for sphere-based applications.

In the network layer, any given communication mechanism may be used, such as TCP/IP, Ethernet, or 802.11. Its management is entirely an internal matter of the contact layer. Hence, the contact layer allows to hide networking issues from the spheres so that sphere-based applications can be written in a clean way. Thus, it remains to specify the spheres layer, the contact layer, and the interface between them.

2.2 The spheres layer

All computation and storage in the spheres layer is organized into so-called *spheres*. A sphere is an atomic thread with its own, private resources that are only accessible by the sphere itself. "Atomic thread" means that a sphere must be completely stored within a single peer and that operations within a sphere are executed in a strictly sequential, non-preemptive way. A prerequisite for this approach to work is that all elementary operations must be strictly non-blocking so that a sphere will never freeze in the middle of a computation. A sphere cannot access any of the resources outside of its private resources. The only way a sphere can interact with the outside world is by sending messages to other spheres. A sphere is bound to the peer that created it. This is important to guarantee some basic properties for the exchange of messages between spheres specified below, but it does not restrict the universality of our approach.

2.3 The contact layer

All communication in the contact layer is handled via so-called *contact points*. Contact points can be thought of as ports but are much more general than that. A contact point is an atomic object that is bound to the sphere that created it. Given a sphere s, let $C(s)$ be the set of contact points created by it, and given a contact point c, let $R(c)$ be the set of spheres that have registered for it. Let *Contact* be the type of a contact point and *FarRef* be the type of a far reference. Contact points are manipulated by the following operations:

- *Contact* **NewContact**(): this creates and returns a new contact point c to the sphere s calling the operation. Formally, $C(s) = C(s) \cup \{c\}$.
- *FarRef* **Reference**(*Contact* c): this returns a far reference to contact point c.
- void **Register**(*Contact* c): this allows a sphere s to register for contact point c. Formally, $R(c) = R(c) \cup \{s\}$.

- void **Unregister**(*Contact* c): this allows a sphere s to unregister for contact point c. Formally, $R(c) = R(c) \setminus \{s\}$.

We demand that **Register** and **Unregister** operations from the same sphere to the same contact point are executed in FIFO order, i.e. in the order they were called, so that eventually $R(c)$ arrives at a correct state.

A contact point c does not need a destructor. As long as $R(c) = \emptyset$, c will not be stored explicitly but only if some sphere registers for it. Variables of type *Contact* or *FarRef* can be passed from one sphere to another. Ownership of a variable of type *Contact* allows to register and unregister for that contact point, and ownership of a variable of type *FarRef* allows to send a message to the contact point. For those familiar with capability-based systems, adding unguessable authentication information to these types turns them into capabilities so that contact points are protected against unauthorized access.

Given a contact point c with far reference r, the "$\leftarrow$" operator allows to send a message m to any sphere that has registered for c. Or more formally, if sphere s calls $\boldsymbol{r \leftarrow m}$, then the contact layer of s first checks whether s has already received a *representative* $s' \in R(c)$ from c. If not, s sends a request to c which returns any $s' \in R(c)$. Afterwards, s forwards the message to s'. If s' unregisters or becomes unavailable, the contact layer of s will request a new sphere from c.

The $\leftarrow$ operation is a *non-blocking, eventual* send operator that guarantees the following properties:

- FIFO ordering: Messages sent by sphere s to some sphere s' via contact point c arrive at s' in the same order in which the $\boldsymbol{r \leftarrow m}$ operations were executed in s.
- At-least-once delivery: Messages are delivered to s' in an at least once fashion. (Notice that exactly-once delivery cannot be guaranteed in a potentially unreliable network.)

2.4 Creating and deleting spheres

Let $S(p)$ be the set of spheres that are currently at peer p. Spheres are created and deleted via the following operations.

- void **Spawn**($C(\cdot)$): this creates a new sphere s in the peer owning the sphere calling **Spawn**, using constructor $C(\cdot)$ for initializing s. Formally, $S(p) = S(p) \cup \{s\}$. Notice that no return value is provided. This makes sure that **Spawn** is non-blocking and the sphere calling **Spawn** has no initial access to s (which is useful for a formal security analysis).
- void **Halt**($D(\cdot)$): this unregisters the sphere s calling **Halt** from all contact points and waits until all tasks have been processed. (Notice that this happens in the background, i.e. the sphere will go on with executing code without waiting for the **Halt** operation to complete.) Afterwards, the destructor $D(\cdot)$ is called in s (which can be used, for example, to save the final state information of s or move it to another peer to recreate s there). Formally, $S(p) = S(p) \setminus \{s\}$.

For security reasons, the **Spawn** operation should only be allowed to be called by a special root sphere r which is created when the peer belonging to r initiates the platform providing the contact layer. In this way, the peer is under the control of which and how many spheres are spawned in it.

Notice that our model only requires seven primitives. All of these primitives are *necessary* in a sense that they cannot be replaced by a combination of the other operations. On the other hand, the primitives are also *sufficient* in a sense that they potentially allow a sphere to reside at any peer in the system and to communicate with any other sphere in the system. Thus, our framework is universal. Furthermore, it satisfies the demands we stated earlier on a distributed programming paradigm. Finally, notice its striking similarity with peer-to-peer systems: spheres are connected by a pointer structure in a similar way as peers

are connected by an overlay network. In fact, the spheres framework allows to implement every peer-to-peer system that has been proposed so far, and much more.

Of course, one may wonder at this point whether the spheres framework will ever be more than just some thought experiment, but interestingly it is already much more than that. There is already a full-grown language support for it, and it allowed to solve an open problem about searchable peer-to-peer systems.

2.5 Actors and the E language

As mentioned above, the basic ideas behind our spheres framework go back to the actors model of Carl Hewitt [12]. However, despite the similarities, the (original) actors model differs in several respects from our spheres model. For example, actors are not required to stay at their location. Instead of using contact points, messages are directly exchanged between actors. Moreover, exactly-once delivery is required, which is impossible to guarantee in an unreliable environment. Though the actors approach has not (yet) made it into a mainstream approach in distributed computing, researchers in the programming language area have continued working on and extending Hewitt's ideas [4, 5, 6] which led, among other results, to the language E [28]. E is a Java-based language that was built from ground up with security in mind. That is, security is not an option but an integral part of E. In fact, E is a capability-based language in which all far references are protected by so-called Swiss numbers. It has advanced network recovery and short-circuiting features that make sure that distributed applications are executed in a robust and efficient way. Actors or spheres are represented by VATs in E, and nothing else is allowed besides VATs and references between them. Also, E provides an advanced form of the "$\leftarrow$" operator. Hence, E follows the philosophy of our approach. However, E does not (yet) provide contact point functionality and at-least-once message delivery guarantees but corresponding extensions are planned in the future. For more information on E, see `http://www.erights.org`. The E language is currently used in the "Theory of Network Communication" course at Johns Hopkins which introduces the students to the state of the art in overlay network concepts. Students use E to implement real peer-to-peer overlay networks in a matter of a week with robustness and security features that would normally take months to implement from scratch. For more information about the course see `http://www.cs.jhu.edu/~scheideler/courses/600.348_F04/`.

2.6 First theoretical results

The spheres model does not only appear to be useful in practice but also in theory. Very recently, it has been used by the second author to compose two peer-to-peer data structures into a single peer-to-peer data structure that combines the advantages of Chord [33] concerning load balancing of data with the advantages of skip graphs [7, 23] concerning the searchability of data. The result is a load-balanced peer-to-peer data structure that can execute insert, delete, and prefix search operations in $O(\log n)$ time, which was previously an open problem [8]. The composition needs certain assumptions for the correctness proof that require a formal framework, and the spheres framework turned out to be sufficient for this.

References

[1] *Message Passing Interface Forum*. `http://www.mpi-forum.org`.

[2] *The Orca Parallel Programming Language*. `http://www.cs.vu.nl/orca/`.

[3] *Split-C*. `http://www.cs.berkeley.edu/projects/parallel/castle/split-c/`.

[4] G. Agha. *Actors: A Model of Concurrent Computation in Distributed Systems*. MIT Press, Cambridge, Massachusetts, 1986.

[5] G. Agha, I. Mason, S. Smith, and C. Talcott. A foundation for actor computation. *Journal of Functional Programming*, 7(1):1–69, 1997.

[6] G. Agha and T. Prasanna. *Formal Methods for Distributed Processing - An Object Oriented Approach*, chapter 8, Actors: A model for reasoning about open distributed systems. Cambridge University Press, 2001.

[7] J. Aspnes and G. Shah. Skip graphs. In *Proc. of the 14th ACM/SIAM Symp. on Discrete Algorithms (SODA)*, pages 384–393, 2003.

[8] B. Awerbuch and C. Scheideler. Peer-to-peer systems for prefix search. Technical report, Johns Hopkins University, Sept. 2004. Available on request. (A preliminary version appeared at Proc. of the 22nd ACM Symp. on Principles of Distributed Computing (PODC), 2003).

[9] Bamboo. `http://bamboo-dht.org`.

[10] H. Bast and T. Hagerup. Fast and reliable parallel hashing. In *Proc. of the 3rd ACM Symp. on Parallel Algorithms and Architectures (SPAA)*, pages 50–61, 1991.

[11] A. D. Birrell and B. J. Nelson. Implementing remote procedure calls. *ACM Transactions on Computer Systems*, 2, 1978.

[12] P. Bishop C. Hewitt and R. Stieger. A universal modular actor formalism for artificial intelligence. In *Proc. of the 1973 International Joint Conference on Artificial Intelligence*, pages 235–246, 1973.

[13] CIRCLE. `http://thecircle.org.au`.

[14] R. Cook. MOD - A language for distributed processing. In *Proc. of the 1st International Conference on Distributed Computing Systems*, 1979.

[15] G. Cugola and G. Picco. PeerWare: Core middleware support for peer-to-peer and mobile systems. Technical report, Politecnico di Milano, May 2001.

[16] F. Dabek, B. Zhao, P. Druschel, J. Kubiatowicz, and I. Stoica. Towards a common API for structured peer-to-peer overlays. In *Proc. of the 2nd International Workshop on Peer-to-Peer Systems (IPTPS)*, 2003.

[17] P. Druschel and A. Rowstron. Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems. In *Proc. of the 18th IFIP/ACM International Conference on Distributed Systems Platforms (Middleware 2001)*, 2001. See also `http://research.microsoft.com/~antr/Pastry`.

[18] P. Eugster, P. Felber, R. Guerraoui, and A.-M. Kermarrec. The many faces of publish/subscribe. To appear in acm computing surveys, Microsoft Research, 2003.

[19] J. Gabarro, C. Martinez, and X. Meseguer. A design of a parallel dictionary using skip lists. *Theoretical Computer Science*, 158(1–2):1–33, 1996.

[20] J. Gil, Y. Matias, and U. Vishkin. Towards a theory of nearly constant time parallel algorithms. In *Proc. of the 32nd IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 698–710, 1991.

[21] GISP. `http://gisp.jxta.org`.

[22] The Object Management Group. *Common Object Request Broker: Architecture and Specification*, omg document number 91.12.1 edition, 1991.

[23] N. J. Harvey, M. B. Jones, S. Saroiu, M. Theimer, and A. Wolman. Skipnet: A scalable overlay network with practical locality properties. In *4th USENIX Symposium on Internet Technologies and Systems*, 2003.

[24] N. C. Hutchinson, L. L. Peterson, M. B. Abott, and S. O'Malley. RPC in the x-Kernel: Evaluating new design techniques. In *Proc. of the 12th Symposium on Operating Systems Principles (SOSP)*, 1989.

[25] M. Jaayeri, C. Ghezzi, D. Hoffman, D. Middleton, and M. Smotherman. CSP/80: A language for communicating sequential processes. In *Proc. of the Distributed Computing CompCon*, 1980.

[26] B. Karp, S. Ratnasamy, S. Rhea, and S. Shenker. Spurring adoption of DHTs with OpenHash, a public DHT service. Technical report, submitted for publication, Carnegie Mellon University, 2003.

[27] MLDonkey. http://www.nongnu.org/mldonkey.

[28] C. Morningstar M.S. Miller and B. Frantz. Capability-based financial instruments. *Financial Cryptography 2000*, 2000. See also http://www.erights.org/elib/capability/ode/.

[29] G. D. Parrington, S. K. Shrivastava, S. M. Wheater, and M. C. Little. The design and implementation of Arjuna. *Computing Systems*, 8(2):255–308, 1995.

[30] W. Paul, U. Vishkin, and H. Wagener. Parallel dictionaries on 2-3 trees. In *Proc. of the 10th International Colloquium on Automata, Languages and Programming (ICALP)*, pages 597–609, 1983.

[31] S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Shenker. A scalable content-addressable network. In *SIGCOMM '01*, 2001.

[32] I. Stoica, D. Adkins, S. Zhuang, S. Shenker, and S. Surana. Internet Indirection Infrastructure. In *Proc. of the ACM SIGCOMM '02*, 2002.

[33] I. Stoica, R. Morris, D. Karger, M.F. Kaashoek, and H. Balakrishnan. Chord: A scalable peer-to-peer lookup service for Internet applications. In *Proc. of the ACM SIGCOMM '01*, 2001. See also http://www.pdos.lcs.mit.edu/chord/.

[34] J. Waldo, G. Wyant, A. Wollrath, and S. Kendall. A note on distributed computing. Technical report SMLI TR-94-29, Sun Microsystems Laboratories, November 1994.

[35] L. Zahn, T. Dineen, P. Leach, E. Martin, N. Mishkin, J. Pato, and G. Wyant. *Network Computing Architecture*. Prentice Hall, 1990.

[36] B.Y. Zhao, J. Kubiatowicz, and A. Joseph. Tapestry: An infrastructure for fault-tolerant wide-area location and routing. Technical report, UCB/CSD-01-1141, University of California at Berkeley, 2001. See also http://www.cs.berkeley.edu/∼ravenben/tapestry.