

OS Support for P2P Programming: a Case for TPS

Sébastien Baehni, Patrick Th. Eugster, Rachid Guerraoui
Swiss Federal Institute of Technology
Distributed Programming Laboratory
IN-R, 1015 Lausanne, Switzerland
{Sebastien.Baehni, Patrick.Eugster, Rachid.Guerraoui}@epfl.ch

Abstract

Just like Remote Procedure Call (RPC) turned out to be a very effective OS abstraction in building client-server applications over LANs, Type-based Publish-Subscribe (TPS) can be viewed as a high-level candidate abstraction for building Peer-to-Peer (P2P) applications over WANs.

This paper relates our preliminary, though positive, experience of implementing and using TPS over JXTA, which can be viewed as the P2P counterpart to sockets. We show that, at least for P2P applications with the Java type model, TPS provides a high-level programming support that ensures type safety and encapsulation, without hampering the decoupled nature of these applications. Furthermore, the loss of flexibility (inherent to the use of any high level abstraction) and the performance overhead, are negligible with respect to the simplicity gained by using TPS.

1. Introduction

Remote Procedure Call (RPC) was first proposed by Birrel and Nelson [1] as a simple abstraction that conceals interactions between remote components beneath traditional procedural interfaces. Partly because of its simplicity and the very facts that it preserves object encapsulation and ensures type safety, and partly because RPC's overhead was very acceptable over sockets, RPC became a dominant paradigm for programming distributed applications over client/server architectures. In these architectures, clients typically communicate with one or several servers following a strongly-coupled request/reply scheme.

With the emergence of Peer-to-Peer (P2P) infrastructures, new forms of decoupled (i.e., anonymous and asynchronous) interactions are needed. One can indeed extend RPC with decoupling flavours. Nevertheless, regardless of

the fact that adding layers over RPC would certainly hamper performance, it is challenging to devise a simple abstraction that could directly fit these architectures and potentially be supported by future Internet-wide operating systems.

So far, P2P developers have generally gravitated towards a few application types: instant messaging (ICQ, AOL's Instant Messenger); collaboration (Aimster, Groove Networks); searching and file sharing (Morpheus, AudioGalaxy); distributed computation (Seti@Home, Parabon). Going beyond these simple applications, and developing more advanced ones, goes through developing basic abstractions for P2P programming.

Some initiatives were recently made towards building libraries or frameworks for deploying P2P applications. A seminal example is the JXTA [2] specification whose implementations provide, for example, protocols for service discovery and many-to-many communication. This specification is rather low level and its protocols can be viewed as the analogous of the basic TCP or UDP protocols [3] for client/server programming over sockets: one needs to explicitly cast types and control encapsulation.

Just like RPC typically hides the underlying mechanisms of sockets and preserves type safety and encapsulation, Type-Based Publish/Subscribe (TPS) [4], a variant of Publish/Subscribe [5], can be viewed as a reasonable candidate abstraction to hide the mechanisms of a low-level P2P library, like JXTA. The distributed event-based interaction scheme promoted by TPS enables the preservation of the decoupled flavor of P2P applications.

This paper presents an implementation of TPS over JXTA and compares the programming and the performance of TPS in writing a typical P2P application with the programming and the performance of using directly JXTA in writing the very same P2P application. Our implementation of TPS over JXTA together with our performance comparisons provide a preliminary, yet interesting, experience towards evaluating the feasibility of equipping fu-

ture Internet-wide operating systems with abstractions like TPS.

Like any high level abstraction, TPS does not apply to all kinds of applications and is obviously less flexible than a lower level library like JXTA. In particular, JXTA simply assumes a common XML knowledge among peers, denoting a very high interoperability, whereas the current implementation of TPS restricts to applications that share the common Java type model. We show that for these applications, the inherent benefits of the use of our TPS library, namely type-safety, encapsulation of application defined event types and code reusability, can be provided without hampering the decoupled nature of P2P computing.

This paper is organized as follows. Section 2 is a brief tutorial on JXTA. Section 3 describes a TPS API and an implementation of this API over JXTA. Section 4 compares the programming of an application using TPS and directly using JXTA (note that the last part of this section has been sketched for limitation purpose, have a look at [6] for the complete version). Section 5 compares the performance of these two implementations. Section 6 summarizes and concludes our experience.

For presentation simplicity and space limitations, we only give excerpts of the interfaces and classes of our implementations. The complete code of our TPS implementation and our testbed applications (both using TPS and directly using JXTA) are available at: <http://lpdwww.epfl.ch> [7]. Please note that this paper relates our experience with the build 29i of JXTA (22nd August). Since then, a lot of changes have been made in the API. We invite you to have a look at our web site to see the latest changes.

2. Background: JXTA

We recall here the basics of JXTA, on top of which we build our TPS abstraction layer. JXTA is a library specification for P2P computing, defining three layers: a core layer, a service layer and an application layer. The application layer wraps all the applications that are developed by JXTA programmers. The service layer is made up of services simplifying the development of the programmer. Various services are currently being implemented by the JXTA community; the best known are the monitoring service, the cms (content management system) service and the wire service (responsible for providing many-to-many communication). The core JXTA layer consists of several protocols ensuring basic communication between the peers, message routing or peer group creation.

2.1. The concepts

The JXTA protocols rely on six concepts: ID, Peer, Pipe, PeerGroup, Advertisement and Message. An ID

identifies any JXTA resource, which can be a peer, a pipe, a peergroup or a codat (code and data). The peer concept points out all networked devices using JXTA. Any device with an electronic pulse could be a JXTA peer (refrigerator, PDA, computer, ...). There are different kinds of peers: "normal" ones and ones that have additional functionalities. Rendez-vous (rdv) are specific peers that keep track of information about peers that are connected. Rendez-vous allow to make the bridge between two different sub-networks. They are mainly used to dispatch information and discovery queries between peers. The second kind of special peers are routers. These are used to route the information from one peer to another if they cannot communicate directly. Peers may have multiple network interfaces.

In order for the peers to communicate, they need a mechanism that does not depend on their network. This mechanism is the pipe. A pipe is a virtual communication channel used to send messages. The basic pipes are asynchronous and uni-directional but some other variants are available (e.g., the very new bidirectional pipes or the many-to-many pipes (called wire)). Pipes are not bound to any physical address (like IP ones). Hence if a peer changes its address, it can continue to use the same pipe for sending or receiving messages. PeerGroups are collections of peers. A peer may join multiple peergroups to share different resources and services. There is no hierarchy inside the groups. A peergroup creates a scoped and monitored environment.

When a new resource (peer, pipe, peergroup, service) is available, a new advertisement is published in order for the other peers to know this resource. An advertisement is a XML message that provides information about the resource. A typical peer advertisement would give information about the network interfaces it provides, about which groups it belongs to, about its name and ID. Each advertisement encompasses an age to distinguish stale advertisements from new ones.

2.2. The protocols

Implementing the JXTA specification consists in implementing the following protocols: Peer Discovery Protocol (PDP, Figure 1), Peer Resolver Protocol (PRP, Figure 2), Peer Information Protocol (PIP, Figure 3), Pipe Binding Protocol (PBP, Figure 4), Endpoint Routing Protocol (ERP, Figure 5) and the RendezVous Protocol (RVP, Figure 6).

The PDP allows different peers to find each other. In fact, this protocol allows to find any kind of published advertisements. Without this protocol, a peer remains alone unless it knows in advance the peers it wants to connect to. This protocol uses the rdv/router peers to improve its performance and the PRP to achieve different discoveries.




Fig. 1: Basic view of the PDP

The PRP is a protocol just above the transport layer. This protocol dispatches each JXTA message to the right services. The more handlers are registered with PRP, the more peers a given peer is potentially able to communicate with.




Fig. 2: Basic view of the PRP

The PIP is used to know the status of a peer. This protocol is responsible for finding and dispatching information about a peer, like the time the peer was up, the different incoming and outgoing channels, the traffic on them, and the different target and source IDs.




Fig. 3: Basic view of the PIP

The PBP is responsible for keeping the different peers of a pipe bound together. Even if the peers are moving in the network (i.e., if their IP addresses do not remain the same), they can continue to use the same pipes to send/receive messages (to achieve that goal, the protocol uses the UUID (Universal Unique Identifier) of the peers).




Fig. 4: Basic view of the PBP

The ERP is used to route the different messages between the different peers. This allows different peers to exchange messages even when they do not know how to connect to each other (because of a firewall for example).




Fig. 5: Basic view of the ERP

Finally, the RVP is used by the rdv peers to dispatch the information (or queries) between the connected peers of the rdv. The RVP uses a cache for performance purpose, implying that a rdv peer must have some additional resources.




Fig. 6: Basic view of the RVP

3. TPS over JXTA

This section overviews the design and implementation of our TPS abstraction over JXTA.

3.1. TPS: Overview

The publish/subscribe paradigm is a communication pattern that provides *time*, *space* and *flow* decoupling among communicating entities. More precisely, the publishers and the subscribers (a) do not need to be up at the same time (time decoupling), (b) do not need to know each other (space decoupling) and (c) the sending/receiving of messages does not block the participants (flow decoupling). This paradigm perfectly suits decoupled networks and serverless architectures. In the original pattern (e.g., [8]), publishers publish information on a subject and subscribers subscribe to subjects. These different subjects are often arranged in hierarchies (specified by a URL-like notation). More advanced communication schemes can be obtained through content-based subscribing, where subscribers express interests in events with particular native properties (e.g., [9]). In our Type-based Publish/Subscribe (TPS) scheme (see Figure 7), the subject is the event object type and the content is the state of instances of that type. Moreover, TPS ensures type safety and preserves event encapsulation with application-defined event types:

the subscriber knows in advance the type of events it receives (type-safety) and subscriptions operations of the type can be used for content-based filtering (encapsulation). So one can easily implement content-based publish/subscribe (hence subject-based) using TPS.




Fig. 7: Type-Based Publish/Subscribe (TPS)

3.2. Generic Java

Our TPS implementation we relate in this paper is based on *genericity*. Using TPS for a specific type T can be viewed as using instances of generic classes with a type parameter instantiated with T. Such generic classes are supported by several languages like C++ (*template*) and Ada (*generic*), while Java supports generics by the idiom of replacing variable types by the top of the type hierarchy. For such languages lacking generic types and methods, adequate extensions have been widely studied. In the case of Java, several solutions have been proposed like Generic Java (GJ) [10] which we have used for our implementation.

Our implementation uses the 1.3 version of the 14th Java Specification Request (JSR), based on GJ, which is expected to be included in the 1.5 version of Java (The 1.3 version of the 14th JSR is a fully Java compatible compiler and enables the use of the original Java Virtual Machine (JVM)).

3.3. The TPS API

The different methods a programmer can use to express a TPS interaction are regrouped within our `TPSInterface`. The corresponding source code is given in Figure 8.

We briefly describe below each of these methods:

(1): This method is used to publish an instance of a type (`Type`) which can be any application-defined type. This instance is sent as an event to the subscribers.

(2): This method (as the next one, i.e., (3)) is used to subscribe to the events of a specific `Type`. Two parameters must be provided: (a) a call-back object which is used to handle received events and (b) a handler for the exceptions that may be raised while handling the received events.

(3): This alternative subscription method is used to register several call-back objects to handle the events in different ways. It is very useful, for instance, if we want to display the complete description of the events in a console and have a sketch of them in a GUI at the same time (for example, see Figure 11).

(4): This method is used to unsubscribe a specified call-back object and its associated exception handler. By doing so, only the specified call-back object is removed.

(5): This method is used to remove all the call-back objects registered so far. After this call, no event is received anymore.

(6, 7): The last two methods are used to obtain the entire set of events received or sent so far.

The other type of the API that the programmer needs to handle is `TPSEngine`. This class gives a reference to the `TPSInterface`. Here is the sketched source code of the `TPSEngine` class:

```
public class TPSEngine<Type> {
    public TPSEngine() {...}
    public TPSInterface
        newInterface(String n, Criteria c,
                    Type t, String[] arg){
        ...
    }
}
```

The programmer uses this class in the initialization phase to get the `TPSInterface`.

```
public interface TPSInterface<Type> {
    public void publish(Type type) throws PSEException; // (1)
    public void subscribe(TPSCallBackInterface<Type> tpsCBI,
                        TPSEExceptionHandler<Type> tpsExH)
                        throws PSEException; // (2)
    public void subscribe(TPSCallBackInterface<Type>[] tpsCBI,
                        TPSEExceptionHandler<Type>[] tpsExH)
                        throws PSEException; // (3)
    public void unsubscribe(TPSCallBackInterface<Type> tpsCBI,
                        TPSEExceptionHandler<Type> tpsExH)
                        throws PSEException; // (4)
    public void unsubscribe(); // (5)
    public Vector objectsReceived(); // (6)
    public Vector objectsSent(); // (7)
}
```

Fig. 8: The TPSInterface

3.4. Architecture

The TPS layer fits between the “standard” application layer and JXTA’s one. In our architecture (Figure 9 and Figure 10), one type is represented by one advertisement. When a subscriber subscribes to a type, it must specify an object implementing the `TPSCallBackInterface` for that type to handle the events and an `TPSEExceptionHandler` (see Section 4.3.3 for an implementation example) responsible for handling the exceptions that may occur while dispatching the events. Our TPS layer is made up of four building blocks (see Figure 10):

- `TPSEngine`: This block is the core of our service. It collects and dispatches the subscriptions and publications.

- Advertisements (Advs): This block is responsible for creating a new advertisement for the type we are interested in as well as for finding and collecting the multiple advertisements that are in relation with our type.
- Interface Repository (IR): This block stores all the call-back interfaces and exception handlers. It also starts and stops the subscriptions
- Connections: This block creates readers, input pipes and output pipes from an advertisement. It sends and receives new messages with the underlying JXTA-WIRE service.

You can see several things that are missing in this typical pub/sub architecture. We do not have a module that do the routing mechanisms [11] and we neither have one for the filtering of events problem [12]. In fact, we rely only on JXTA for doing the routing of events (and it is one of our main goal, to test what JXTA can do for us). We do not take into account the filtering problem as we want to remain concise in building the first blocks of our TPS architecture (moreover, this problem has been studied in [4]).

Do not considering those aspects can lead to severe performance leaks and problems, especially for large-scale distributed systems and that is definitely what we will implement next in our TPS API.




Fig. 9: General architecture

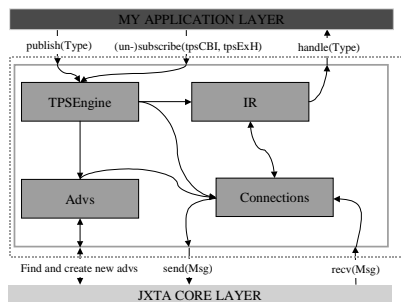


Fig. 10: General architecture (details of the TPS layer)

4. The programming experience

We compare here the programming of a typical P2P application using our TPS abstraction with the programming of the very same application using JXTA directly.

4.1. Ski-Rental Application

If you want to go skiing, you need skis. If you do not have any, you have two possibilities, either you buy them or you rent them. In the latter case, you will typically go to different shops in order to see what kind of skis you want and also to compare the different prices. Of course, nowadays, you could also do that online, by visiting different web-sites. However, you must spend time doing that: you must stay behind your computer trying to find the best skis. A more comfortable way to do that is to use the TPS paradigm over a P2P infrastructure. You would then subscribe to the ski-rental type and wait for the answers. The infrastructure will be responsible for sending the subscription to the other peers and also getting the responses. You can now do something else during the search phase of the program and come back later to get the answers when they are available. Figure 11 depict two GUIs that a publisher and a subscriber have with our ski-rental application.

When the publisher (a shop for example) starts, a search for a SkiRental advertisement is first launched. If the application does not find such advertisement in a specific amount of time, it creates its own one, but keeps trying to find others in order to send messages to the maximum number of interested subscribers. After that, an output pipe is created to send messages, and the window is displayed (see Figure 11).

In Figure 11, we can see that the shop seller can set the different options for the kinds of skis he wants to offer for rent (the brand, the duration of the rental, the price, ...). After setting these parameters, he just has to click on the publish button and the proposition is sent to all the interested subscribers via the output wire pipe.

For the subscriber, the same kind of initialization is done as for the publisher. Once this initialization has been accomplished, the window is displayed. In this interface, the subscriber can see the different propositions from publishers of ski-rental advertisements. After some time, the subscriber can choose the best propositions and, maybe, send a message to the shop (in hitting the answer button).



Fig. 11: GUI of a ski-rental publisher and subscriber

4.2. Programming Phases

Programming a TPS application, like the ski-rental, can be divided into four main phases, as depicted in Figure 12 (the arrows conveys the causality). In our implementation, we associate one instance of a publish/subscribe engine per type. If a publisher (or a subscriber) is interested in

several “unrelated” types (i.e., different types that do not belong to the same type hierarchy), several instances of the publish/subscribe engine for each type of interest must be created (leading to a loss of resources, but some optimisations can be made to reuse the modules of a former pub/sub engine). In Figure 12, this is conveyed by the fact that the type definition phase precedes the initialization phase.




Fig. 12: The 4 different phases

In the following, we overview the two different ways of developing the ski-rental application according to the four phases, first using our TPS API and second using directly JXTA.

4.3. Renting skis with TPS

We present here the different phases shown in Figure 12 to create a simple application using our TPS architecture.

4.3.1. Type definition phase: the `SkiRental` type. We give here the basic type used in our application. This type contains the name of the renter, the price, the brand of the skis and the number of days the skis need to be rented. Here is the sketched source code of our simple type:

```

public class SkiRental implements
    Serializable {
    public SkiRental(String shop, float price,
        String brand, float days)
        {...}
}
  
```

4.3.2. Initialization phase. When a user wants to use the TPS API, he must first write few lines to initialize the publish/subscribe engine:

```

TPSEngine<SkiRental> tpse =
    new TPSEngine<SkiRental>();
TPSInterface tpsInt =
    tpse.newInterface("JXTA", null,
        new SkiRental(), argv);
  
```

In the first line, we create the publish/subscribe engine and specify the type of interest. In the second line, the second parameter specifies a criteria we want for filtering advertisements (may be null). The third parameter is an instance of the type of the events we are interested in. We must provide this instance because GJ does not provide

runtime information about (actual) type parameters. The last parameter denotes the arguments of the main class (may be null).

4.3.3. Subscription phase. To subscribe to events, one must create two objects (as described when presenting the `subscribe()` method, Figure 8): one implementing the `TPSCallBackInterface` interface and another one implementing the `TPSExceptionHandler` interface. Here is an implementation of the first interface:

```

public class MyCBInterface implements
    TPSCallBackInterface<SkiRental> {
    public void handle(SkiRental skiR)
        throws CallBackException {
        System.out.println(skiR);
    }
}
  
```

This class defines what needs to be done when new events are received. In this case, we just print the events into the console. Here is a sketched implementation for the second interface:

```

public class MyExHandler implements
    TPSExceptionHandler<SkiRental> {
    public void handle(Throwable th) {...}
}
  
```

Besides these two classes, here are the lines one must add to subscribe to the type (`SkiRental`) specified before:

```

MyCBInterface mCBInt = new MyCBInterface();
MyExHandler mExH = new MyExHandler();
tpsInt.subscribe(mCBInt, mExH);
  
```

4.3.4. Publication phase. Up to now, we have only seen the subscriber perspective. If a publisher wants to publish an instance of the `SkiRental` type, here is the line he must add after the initialization phase:

```

tpsInt.publish(
    new SkiRental("XTremShop", "Salomon",
        14f, 100f);
  
```

4.4. Renting skis with JXTA

Our aim here is to create the very same application than the one with TPS, i.e., an application with the same functionalities¹ as TPS. Unfortunately, for limitation purpose, we are not able to give you the class files mandatory for developing the architecture presented in Figure 9, Figure 10.

These classes should have convinced you that TPS hides a lot of programming details. For example, writing the very same application with JXTA implies writing, at

1. (1) Minimization of the number of advertisements for the same type, (2) management of multiple advertisements at the same time and (3) handling of duplicate messages.

least, about 4000 lines of code more than using directly TPS. Moreover, TPS allows the programmer to focus only on the portion of the code he is interested in. This prevents from spending time learning the underlying JXTA concepts. Finally, TPS prevents the programmer from performing wrong type casts at runtime, and hence saves precious debugging time. If you want to see this classes, either have a look at [6] or at [7].

5. The performance experience

This section presents the performance results of both our ski-rental application based on our TPS layer (SR-TPS), and implemented directly with JXTA (SR-JXTA). Even if JXTA-WIRE alone is not comparable with SR-TPS and SR-JXTA (since it does not insure the properties described in Section 4.4), we use it here as a (lower bound) reference point.

We used the following computer configurations: Sun Ultra 10 (CPU 440 MHz, RAM 256 MB) on Solaris 7; FastEthernet (100 Mbits/s); JXTA version 1.0 (build 29i, 08-22-2001); Java version "1.4.0-beta" (Java(TM) 2 Runtime Environment, Standard Edition (build 1.4.0-beta-b65), Java HotSpot(TM) Client VM (build 1.4.0-beta-b65, mixed mode)); messages size: 1910 bytes.

We give here the throughput for a limited number of participants (at the time of our implementation, JXTA was not able to handle connections between more than 5 peers sending a lot of messages). These tests aim at giving a hint about the differences between the three implementations. Again for a more complete set of results, see [6].

5.1. Throughput: the publisher viewpoint




Fig. 13: Publisher's throughput

We consider here a set of 100 published events and we measure the time for the publisher to deliver those events to the subscriber(s). Again, the values for SR-JXTA and SR-TPS are very close. We can also notice that our differ-

ent layers are slightly slower than JXTA-WIRE itself (e.g., about two events per second for one subscriber) (see Figure 13). When the number of subscribers increases, the differences between the layers become insignificant (e.g., with four connected subscribers, only 0.3 events per second between JXTA-WIRE and SR-JXTA and 0.5 events per second between JXTA-WIRE and SR-TPS).

5.2. Throughput: the subscriber viewpoint

Here the publishers try to flood the subscriber (10000 events published per each publisher). Every second, we measure the number of events that are received; during 50 seconds. The results are given in Figure 14. Once again, we have a quite big standard deviation and the number of events received per second is not really stable. For example, with a single publisher, the average throughput for JXTA-WIRE is about 7.8 events per second and, for SR-JXTA and SR-TPS, the values are 6.1 and 6.0 respectively. If we compare these results with the ones from Figure 13, we can see that, for one publisher, JXTA-WIRE saturates. JXTA-WIRE can simply not handle all published events (e.g., about nine per second for one subscriber (see Figure 13)).

When we increase the number of publishers, the average number of events received per second remains quite the same for the different layers. Again, the average drops by a factor of about three (because the subscriber must handle more connections (as there are more publishers)).




Fig. 14: Subscriber's throughput

6. Concluding Remarks

This experience paper makes a case for TPS (Type-Based Publish-Subscribe) as a viable abstraction for future Internet-wide operating systems to support P2P applications. TPS fits particularly well the decoupled nature of server-less P2P applications. In short, TPS is simple to use, ensures type-safety and encapsulation and yet preserves the decoupled nature of P2P applications. This pa-

per describes a TPS API and an implementation of TPS over JXTA, and then compares the programming and performance of a testbed application over TPS and directly over JXTA. Roughly speaking, (1) TPS makes the programming of a P2P application significantly easier than using directly a library like JXTA and (2) does not introduce a significant overhead with respect to JXTA.

Our current TPS prototype is based on the JXTA release of August 22, 2001. New implementations of JXTA will obviously impact our prototype but we do not believe they will fundamentally impact the nature of the results drawn from the present experience.

Of course, more programming and implementation testbeds need to be performed before TPS can be realistically viewed as a reasonable general abstraction for P2P applications. In particular, measuring the lack of programming flexibility that our abstraction involves is not clear. We can for example easily see through our ski-rental application that our TPS API does not enable a subscriber to immediately reply to a publisher that posted an interesting event. This would require a combination with a more traditional RPC kind of interaction or directly using the underlying P2P library. Another loss of flexibility is our assumption that the different peers must a priori agree on the Java type system which is not the case when using JXTA directly. Figuring out “loose” ways of achieving such common knowledge at run-time (e.g., by representing types through XML data structures) is the subject of ongoing investigations.

Finally, we are thinking of a new architecture allowing us to use your own network and routing protocols improving the reliability and performance of TPS. One possibility would be to implement a hierarchical probabilistic multicast algorithm [4] between a low-level network protocol (like UDP) and our high-level TPS abstraction.

References

- [1] Andrew D. Birrell and Bruce Jay Nelson. Implementing Remote Procedure Calls. In *Proceedings of the 9th ACM Symposium on Operating Systems Principles (SOSP'83)*. October 1983.
- [2] Project JXTA web site. <http://www.jxta.org>. Sun Microsystems. 2001.
- [3] A. S. Tanenbaum. *Computer Networks*. Prentice-Hall, third edition. January 1996.
- [4] P. Th. Eugster. Type-based publish/subscribe. *PhD thesis*. Swiss Federal Institute of Technology, Lausanne. December 2001.
- [5] B. Oki, M. Pfluegl, A. Siegel et al. The Information Bus - An Architecture for Extensible Distributed Systems. In *Proceedings of the 14th ACM Symposium on Operating System Principles (SOSP '93)*. December 1993.
- [6] S. Baehni, P. Th. Eugster, and R. Guerraoui. OS Support for P2P Programming: a Case for TPS. *Technical Report*. February 2002.
- [7] LPD web site. <http://lpdwww.epfl.ch>. For the full source codes, follow People / Sebastien Baehni / Current work.
- [8] TIBCO. *TIB/Rendezvous White Paper*. <http://www.rv.tibco.com>. 1999.
- [9] *Gryphon: Publish/Subscribe over public networks*. IBM T.J. Watson Research Center. <http://researchweb.watson.ibm.com/gryphon/Gryphon/gryphon.html>. February 2001.
- [10] G. Bracha, M. Odersky, D. Stoutamire and Ph. Wadler. Making the future safe for the past: Adding genericity to the Java programming language. In *Proceedings of the 13th ACM Conference on Object-Oriented Programming Systems, Languages and Applications (OOPSLA '98)*, pages 183-200. October 1998.
- [11] A. Rowstron and P. Druschel. Pastry: Scalable, Decentralized Object Location, and Routing for Large-Scale Peer-to-Peer Systems. In *Proceedings of the 3rd Middleware conference (Middleware 2001)*. November 2001.
- [12] M. Aguilera, R. Strom, D. Sturman, M. Astley and T. Chandra. Matching events in a content-based subscription system. In *Proceedings of the 2nd ACM Special Interest Group on Management Of Data (SIGMOD'99)*. May 1999.
- [13] A. Carzaniga, D.S. Rosenblum and A.L. Wolf. Design and evaluation of a wide-area event notification service. *ACM Trans. on Computer Systems*, 19(3):332-383. August 2001.
- [14] I. Clarke, O. Sandberg, B. Wiley et al. Freenet: A Distributed Anonymous Information Storage and Retrieval System. In *Proceedings of the International Computer Scientist Institute Workshop on Design Issues in Anonymity and Unobservability (ICSI'2000)*. July 2000.
- [15] A. Carzaniga, D. S. Rosenblum and A. L. Wolf. Achieving Scalability and Expressiveness in an Internet-Scale Event Notification Service. In *Proceedings of the 19th ACM Symposium on Principles of Distributed Computing (PODC'00)*. July 2000.
- [16] D. Heimbigner. Adapting Publish/Subscribe Middleware to Achieve Gnutella-like Functionality. In *Proceedings of the 16th ACM Symposium on Applied Computing (SAC'2001)*, pages 176-181. 2001.
- [17] John A. Zinky, Linsey O'Brien, David E. Bakken, Vijaykumar Krishnaswamy, Mustaque Ahamad. Pass-A service for Efficient Large Scale Dissemination of Time Varying Data Using CORBA. In *Proceedings of the 19th International Conference on Distributed Computing Systems (ICDCS'99)*. May 1999.