

Parallel-Distributed Programming Tool for the Linda Programming Model

Ernest Scheiber

Abstract

This note presents an implementation of the Piranha model to program parallel - distributed application and a development tool. We use the TSpaces framework from I.B.M. as a middleware for data exchanges. Our implementation does not introduce a new API but requires certain program structure.

1 Introduction

Even if there is a lack of agreements on what exactly grid computing or cloud computing are, the parallel - distributed programming API's or tools are one of the main ingredient to build such computing environment or application.

Linda is not a programming model but a set of operations that can be added to any language. Linda gives a way to communicate and synchronize different processes. We know the following Linda type development frameworks in Java: *TSpaces* from I.B.M., [6] and *JavaSpaces* from SunMicrosystems.

The *Piranha* model is a parallel - distributed computing model based on Linda, [2]. The paradigm behind Piranha is that of dispatcher-worker. The dispatcher coordinates the activities while the workers solve the received tasks. Between them, the communications are programmed using the Linda operations.

Other known parallel - distributed computing models are MPI, [5, 1], map-reduce.

This note presents an implementation of the Piranha model to program parallel - distributed application and a development tool. We use the TSpaces framework, but instead other messaging service may be used. The actual messaging services (*Sun MessageQueue*, *apache-activemq* or *qpuid*) own the required functionality of a Linda framework. The reason of our development activities consists in the lack of free correspondent software. A preliminary version was reported in [3].

An application to be used with the tool is constrained to a specific structure; it is composed from a dispatcher thread and worker threads. Between the dispatcher thread and the worker threads there are asynchronous message changes. These messages are kept by a TSpaces server until they are consumed. The tool allows us to state the network of workstations involved in the computation, to deploy and to launch an application.

2 The Structure of an Application

We suppose that several computers in a network will perform the required computation. The Java application contains:

- A *Console* class - The Console launches to run the *Dispatcher*.

- On any workstation there is a *Service* class, it instantiates and starts the *Workers*.
- The application associated TSpace server, it keeps the messages until they are consumed.

The structure of an application is outlined in Fig. 1.

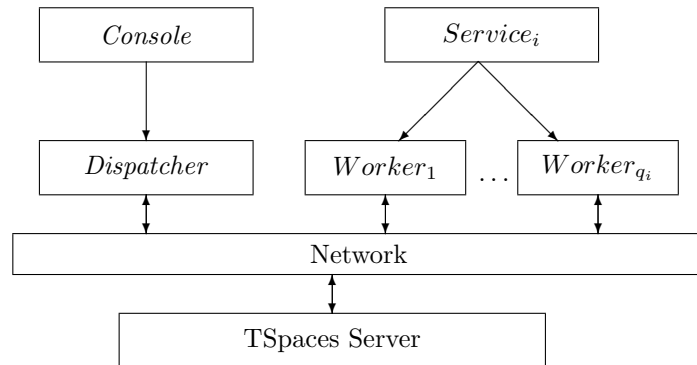


Figure 1: The structure of an application.

The *Dispatcher* and the *Worker* classes are threads. The constructors of the dispatcher and of the worker classes are responsible to establish the connection with the TSpaces server. The run methods of these Java threads contain the specific activities to solve the given problem. On a workstation it is possible to start several worker threads. The i index in Fig. 1 denotes a workstation on which q_i worker threads are run. The *Console* and *Service* classes are independent from an application.

Between the dispatcher thread and the worker threads there are asynchronous message exchanges. These messages are kept by a TSpaces server. The data to be exchanged between the dispatcher and the workers are wrapped into objects. To distinguish between different kinds of data, a tag field may be introduced. A message consumer (dispatcher or worker) waits until all the required messages are available. In this way the synchronization problems are solved as well as the coordination of the activities.

3 The anatomy of our implementation

We shall analyze the codes of the above mentioned classes for a very simple application: the dispatcher sends a greeting message to each worker. After receiving the message, a worker responds indicating the name of its computer.

Let us suppose that we use n workstations and that on the i -th station there will be launched q_i workers ($\sum_{i=1}^n q_i = q$).

Some environment data are fixed in two property files:

propConsole		propService	
Name		Name	
AppName		AppName	
Host		Host	
Port		Port	
TasksNumber	q	TasksNumber	q
		ComputerTaskNumber	q_i
		ComputerFirstTaskIndex	$\sum_{j=1}^{i-1} q_j, q_0 = 0$

where *AppName*, *Host*, *Port* identifies the name of the application TupleSpace and respectively the host of the *TSpaces* server with the corresponding port number.

As we have mention, the dispatcher is launched by the *Console* class, while the workers are launched by the *Service* class. These two classes are independent of an application.

The <i>Console</i> class	The <i>Service</i> class
<pre> package *****; import java.util.Properties; import java.io.*; public class Console{ Properties p=null; public Console(String path2Prop) { try{ // Loads the properties FileInputStream fis=new FileInputStream(path2Prop+"propConsole"); p=new Properties(); p.load(fis); } catch(Exception e){ System.out.println(e.getMessage()); System.exit(1); } } public static void main(String[] args){ String fs=System.getProperties(). getProperty("file.separator"); String path2Prop; // An argument may contain the path to the property file if(args.length>0) path2Prop=args[0]+fs; else path2Prop=""; Console obj=new Console(path2Prop); // Set the properties String appName=obj.p.getProperty("AppName"); String host=obj.p.getProperty("Host"); String sPort=obj.p.getProperty("Port"); String sTasks=obj.p.getProperty("TasksNumber"); int port=Integer.parseInt(sPort); int tasks=Integer.parseInt(sTasks); // The dispatcher is launched into execution Dispatcher dispatcher=new Dispatcher(appName, host,port,tasks); dispatcher.start(); } } </pre>	<pre> package *****; import java.util.Properties; import java.io.*; public class Service{ Properties p=null; public Service(String path2Prop) try{ // Loads the properties FileInputStream fis=new FileInputStream(path2Prop+"propConsole"); p=new Properties(); p.load(fis); } catch(Exception e){ System.out.println(e.getMessage()); System.exit(1); } } public static void main(String[] args){ String fs=System.getProperties(). getProperty("file.separator"); String path2Prop; if(args.length>0) path2Prop=args[0]+fs; else path2Prop=""; Service obj=new Service(path2Prop) String appName=obj.p.getProperty("AppName"); String host=obj.p.getProperty("Host"); String sPort=obj.p.getProperty("Port"); String sTasks=obj.p.getProperty("TasksNumber"); String sIndex=obj.p.getProperty("ComputerFirstTaskIndex"); String sComputerTasks=obj.p.getProperty("ComputerTaskNumber"); int port=Integer.parseInt(sPort); int tasks=Integer.parseInt(sTasks); int index=Integer.parseInt(sIndex); int computerTasks=Integer.parseInt(sComputerTasks); // The required number of workers are // launched into execution Worker[] worker=new Worker[computerTasks]; for(int i=0;i<computerTasks;i++){ worker[i]=new Worker(appName,host, port,tasks,index+i); worker[i].start(); } } } </pre>

The wrapper class that contain the data to be exchanged between the dispatcher and workers may be

```

1 package tshello;
2 import java.io.Serializable;
3
4 public class DataWrapper implements Serializable{
5     // tag=0 message sent by the dispatcher
6     // tag=1 message sent by a worker
7     public int tag;
8     public String mesaj;
9
10    public DataWrapper(String mesaj, int tag) {
11        this.mesaj=mesaj;
12        this.tag=tag;
13    }
14 }

```

The connection to the TupleSpace server is made in the constructors of the *Worker* and *Dispatcher* classes.

The code of the *Worker* class is

```

1 package tshello;
2 import com.ibm.tspaces.*;
3 import java.net.*;
4 import java.io.*;

6 public class Worker extends Thread{
7     private int id;
8         private TupleSpace ts=null;
9         private String tsName;
10    private int tasks;
11    private PrintStream f;

13    public Worker(String tsName, String host, int port, int tasks, int id){
14        this.tsName=tsName;
15        this.id=id;
16        this.tasks=tasks;
17        ts=startTupleSpace(tsName, host, port);
18        try{
19            f=new PrintStream("worker"+id+".txt");
20        }
21        catch(Exception e){
22            System.out.println("File error :"+e.getMessage());
23        }
24    }

26    public void run(){
27        Tuple tuple=null;
28        DataWrapper d=null;
29        String source=tsName+id;
30        String dest=tsName+" "+id;
31        try{
32            Tuple template=new Tuple(source, new FieldPS(DataWrapper.class));
33            tuple=(Tuple) ts.waitToTake(template);
34            d=(DataWrapper) tuple.getField(1).getValue();
35            if(d.tag==0){
36                String msg="Message received by "+id+" from the dispatcher:\n"+d.mesaj;
37                System.out.println(msg);
38                f.println(msg);
39                InetAddress addr=InetAddress.getLocalHost();
40                String s=addr.getHostName();
41                String mesOut="Hello from "+id+" at "+s;
42                d=new DataWrapper(mesOut,1);
43                FieldPS ps=new FieldPS(d);
44                ts.write(dest, ps);
45                f.close();
46            }
47        }
48        catch(Exception e){
49            System.out.println("Exception-Worker : "+e.getMessage());
50        }
51    }

53    private TupleSpace startTupleSpace(String tsName, String host, int port){
54        TupleSpace ts=null;
55        try{
56            Tuple active=TupleSpace.status(host, port);
57            if((active==null)|| (active.getField(0).getValue().equals("NotRunning"))){
58                System.out.println("TupleSpace Server is not available");
59                System.exit(1);
60            }
61            ts=new TupleSpace(tsName, host, port);
62        }
63        catch(TupleSpaceException e){
64            System.out.println("TupleSpaceException "+e.getMessage());
65            System.exit(1);
66        }
67        return ts;
68    }
69 }

```

Finally, in the *Dispatcher* class, there are send messages to the workers (the *scatter* method) and are waiting to receive the corresponding responses (the *gather* method):

```

1 package tshello;
2 import com.ibm.tspaces.*;
3 import java.io.*;

5 public class Dispatcher extends Thread{
6     private int tasks;
7     private TupleSpace ts=null;
8     private String tsName;
9     private PrintStream f;

11    public Dispatcher(String tsName,String host,int port,int tasks) {
12        this.tasks=tasks;
13        this.tsName=tsName;
14        ts=startTupleSpace(tsName,host, port);
15        try{
16            f=new PrintStream("dispatcher.txt");
17        }
18        catch(Exception e){
19            System.out.println("File error :"+e.getMessage());
20        }
21    }

23    public void run(){
24        try{
25            scatter();
26            gather();
27            ts.cleanup();
28        }
29        catch(Exception e){
30            System.out.println(e.getMessage());
31        }
32    }

34    private void scatter(){
35        String mesOut="Hello from the dispatcher !";
36        DataWrapper d=new DataWrapper(mesOut,0);
37        try{
38            FieldPS ps=new FieldPS(d);
39            Tuple multi=new Tuple();
40            for(int i=0;i<tasks;i++){
41                String dest=tsName+i;
42                Tuple nextTuple=new Tuple(dest,ps);
43                multi.add(new Field(nextTuple));
44            }
45            TupleID [] ids=ts.multiWrite(multi);
46        }
47        catch(TupleSpaceException e){
48            System.out.println("TupleSpaceException-scatter0 "+e.getMessage());
49        }
50        System.out.println("Scatter OK");
51    }

53    private void gather(){
54        Tuple tuple,template;
55        DataWrapper d=null;
56        try{
57            for(int i=0;i<tasks;i++){
58                String source=tsName+" "+i;
59                template=new Tuple(source,new FieldPS(DataWrapper.class));
60                tuple=(Tuple)ts.waitToTake(template);
61                d=(DataWrapper)tuple.getField(1).getValue();
62                if(d.tag==1){
63                    System.out.println(d.mesaj);
64                    f.println(d.mesaj);
65                }
66            }
67            f.close();
68        }
69        catch(Exception e){
70            System.out.println("Exception-gather "+e.getMessage());
71        }
72    }

74    private TupleSpace startTupleSpace(String tsName, String host,int port){. . .}
75 }

```

4 An example

The sixteen grid problem (www.ams.org/ams/16-grid.html). Each of the numbers 1,2,...,16 is used exactly once in the empty cells to form arithmetic expressions connected by symbols for the four basic operations. Each row (column) is an arithmetic expression, read and performed left to right (top to bottom), disregarding the usual order of operations, to yield the result at the right (bottom).

	-		×		×		=-60
×	█	+	█	×	█	÷	
	+		+		-		=29
+	█	+	█	-	█	×	
	-		-		-		=-14
+	█	-	█	+	█	-	
	×		-		-		=32
=28		=1		=57		=27	

A solution, presented in [4] is based on a parallelization scheme of the backtracking algorithm. Based on the domain decomposition method, there is a more efficient way to solve this problem. In the set on n order permutations, \mathcal{P}_n , considering the lexicographic order the first permutation is $(1, 2, \dots, n)$ and the last is $(n, n - 1, \dots, 1)$. Relative to this order, it may compute the m -th permutation, $1 \leq m \leq n!$.

The first $q_1 = (n - 1)!$ elements of \mathcal{P}_n have 1 on the first position, the next q_1 elements have 2 on the first position and so on. In the group having i on the first position, there are $n - 1$ subgroups with $q_2 = (n - 2)!$ elements with the second element equals respectively with $1, 2, \dots, i - 1, i + 1, \dots, n$. If the first $s - 1$ elements are fixed, i_1, \dots, i_{s-1} then there exists $n - s + 1$ groups with $q_s = (n - s)!$ elements, having on the i -th position an element of $\{1, \dots, n\} \setminus \{i_1, \dots, i_{s-1}\}$.

To find the s -th element of the m -th permutation, we compute the group number t_s having the i_s element as constant, $t_s = \lceil \frac{m}{q_s} \rceil$, and finally i_s will be the l_s -th element of the sequence $1, \dots, n$, after we have deleted i_1, \dots, i_{s-1} , where $l_s \equiv t_s \pmod{n - s + 1}$, $l_s \in \{1, 2, \dots, n - s + 1\}$.

If there are p workers, then the set of all $n!$ permutations are equally divided into p intervals. Each worker generates sequentially the permutations of its attached interval and verifies the restrictions of the problem.

The unique solution of the sixteen grid problem is

1	-	2	×	4	×	15	=-60
×	█	+	█	×	█	÷	
12	+	8	+	14	-	5	=29
+	█	+	█	-	█	×	
13	-	7	-	9	-	11	=-14
+	█	-	█	+	█	-	
3	×	16	-	10	-	6	=32
= 28		= 1		= 57		= 27	

5 The Development Tool

To run the application a lot of activities are to be performed: to deploy the application, to launch the Service program on each workstation and finally to launch the Console on the host station.

To assist these activities, a development tool was created. The *Service* program becomes a servlet and we use the *apache-tomcat* as a servlet container Web server. The *Console* program which starts the dispatcher makes the requests to the *Service* servlets, too.

To the network of workstations it is associated a second TSpaces server to keep the names of the involved computers - denoted as the tuple space of the network.

The tool contains two parts:

- The *Advertiser* serves to state the network of workstations. Using the *Advertiser* a computer is linked to the network of workstations to perform a parallel-distributed computing.

The graphical interface of the *Advertiser* allows to

- Declare the computer as a member of the network of workstations - i.e. a tuple with the name of the computer is written into the tuple space of the network;
- Remove the computer from the network - i.e. remove the above define tuple from the tuple space of the network;
- Show the list of the computers in the network of workstations.

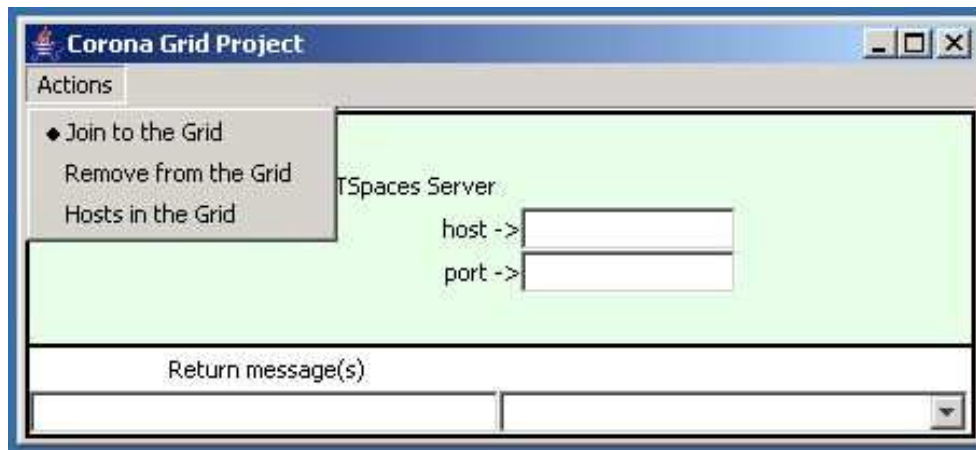


Figure 2: The window of the *Advertiser* tool.

- The *Developer* allows to
 - Generate the folders of the application;
 - Compile and archive the service part of the application;
 - Deploy the service part to the workstations of the network. The deployment is done with the *apache-tomcat-deployer*;
 - Compile the console part of the application.
 - Launch the console to run. In the *Console* class the request are programmed using the *apache commons-httpclient* software;
 - Undeploy the service part of the application.

These targets are executed through *apache-ant*.

The *Developer* is installed only on the host workstation.

A number of parameters are required: The name of the application, The host and the port of the application tuple space, The number of the worker threads, The path to the *apache-ant* and The username and password for the *manager* application of *apache-tomcat*.

The list of the computers name in the network is required, too. This list may be generated with the *Advertiser*, contained in the *Developer*, too.

The scenario to run a parallel-distributed application using the framework involves:

1. To set up the network of workstations:
 - (a) On launches the tuple space associated to the network;

- (b) On each workstation starts the *apache-tomcat* Web server and, with the Advertiser, declares the availability of that computer to join to the network of workstations.
- 2. Using the Developer, on the host workstation, we can install, deploy, start the tuple space associated to the application and execute the parallel-distributed computation launching the Console.

The current version of the framework may be downloaded from the author's Web page <http://cs.unitbv.ro/site/papers/scheiber>. The archive contains several other examples: a numerical integration, quicksort, the queens' problem, the graphic representation of the Mandelbrot set, the dining philosophers' problem.

Conclusions. An implementation of the Piranha model for a parallel-distributed application is developed. The implementation does not introduce a new API but requires certain program structure. In addition we have created a tool to assist the development of an application that can be ported to any Java enabling platform. As a drawback of our approach is that, if a workstation dies then a running application never finishes. We intend to work on this problem. The security constraints are that of the *apache-tomcat* Web server.

References

- [1] R. BISSELING, *Parallel Scientific Computation. A structured approach using BSP and MPI*, Oxford Univ. Press, 2004.
- [2] G. A. PAPADOPOULOS, F. ARBAB, Coordination models and languages. CWI Report, SEN-R9834, 1998.
- [3] SCHEIBER E., 2007, A TSpaces Based Framework for Parallel - Distributed Applications. *Knowledge Engineering Principles and Techniques 1* (2007), Cluj University Press, 341-345.
- [4] SCHEIBER E., 2009, A Parallelization scheme of Some Algorithms. *Knowledge Engineering Principles and Techniques 2* (2009), *Studia Universitatis Babeş-Bolyai, Informatica*, 244-248.
- [5] M. SNIR, D. OTTO, S. HUSS-LEDERMAN, D. WALKER, J. DONGARRA, *MPI: The Complete Reference*. MIT Press, Cambridge, MA, 1996.
- [6] * * * , TSpaces-User's Guide & Programmer's Guide, Distributed with the software, Version 2.1.2, 2000.

Ernest Scheiber
 Transilvania University of Braşov
 Faculty of Mathematics and Computer Science
 str. Iuliu Maniu 50
 ROMANIA
 E-mail: scheiber@unitbv.ro