

Control and Debugging of Distributed Programs Using Fiddle

João Lourenço^{*,1}, José C. Cunha*,
Vitor Moreira*

** Departamento de Informática
Faculdade de Ciências e Tecnologia
Universidade Nova de Lisboa
Portugal*

ABSTRACT

The main goal of **Fiddle**, a distributed debugging engine, is to provide a flexible platform for developing debugging tools. **Fiddle** provides a layered set of interfaces with a minimal set of debugging functionalities, for the inspection and control of distributed and multi-threaded applications.

This paper illustrates how **Fiddle** is used to support integrated testing and debugging. The approach described is based on a tool, called **Deipa**, that interprets sequences of commands read from an input file, generated by an independent testing tool. **Deipa** acts as a **Fiddle** client, in order to enforce specific execution paths in a distributed PVM program. Other **Fiddle** clients may be used along with **Deipa** for the fine debugging at process level. **Fiddle** and **Deipa** functionalities and architectures are described, and a working example shows a step-by-step application of these tools.

KEYWORDS: Distributed Debugging, Software Testing, Tool Integration

1 Introduction

Developing parallel and distributed applications is still a difficult task, even after about several decades of intense research on methodologies and support tools. On one hand, this is due to the intrinsic complexity of distributed computing, involving many dynamic interacting entities, executing on multiple processors, leading to a global nondeterministic behavior. On the other hand, much research and development in this area has been facing a need to continuously adapt to new hardware, operating system and middleware platforms, and new programming language approaches (ranging from imperative to functional or logical, and to object- and agent-oriented models). Such constant evolution in the computing models and platforms made many interesting tools obsolete, as they were not able to adapt, or their design was too dependent upon specific models or platforms.

As parallel machines were becoming available since the 80s, parallel debuggers have been developed and some of them evolved into significant commercial tools, such as **TotalView** [Too]. Initially, such debuggers usually followed a monolithic approach, where a debugging engine and its user interface were combined into a single large program.

Developments on distributed computing have motivated debugging architectures based on the client/server model, such as **p2d2** [Hoo96], which clearly separates the user interface from the debugging engine. Most designs are based on having a separate process in each machine node, to support the remote access to the debugging functionalities, and rely upon a central server to do all

¹E-mail: {jml, jcc, vrm}@di.fct.unl.pt

the processing. This was an important design decision to enable the debugging tool to adapt to different user requirements and environments but the importance of such a step was not recognized until recently. Related and significant efforts were also done concerning monitoring of parallel and distributed programs [CLV⁺98, CD98, LWSB97, LWO96, KP00, TSLW00, KGV96].

An important requirement for flexible development tools is their ability to adapt to new computing platforms or user requirements. This motivates an approach based on a clear separation between the software layers which support the minimal functionalities required by a tool, and the layers which provide the extensions that may be required by specific user and application scenarios. This requires a debugging support platform that enables the inclusion of new tools, to provide complementary functionalities for application development and acting as intermediaries between the debugger and the other tools.

In previous work, some of the above requirements, involving the integration of a testing tool for parallel programs (STEPS) and a parallel debugger (DDBG) were tested [LCK⁺97]. However, the limitations of such experiment have motivated the continuation of this work towards a more flexible debugging architecture, Fiddle.

In this paper, we first describe the main characteristics of the Fiddle architecture. Next, in Sec. 3, we discuss an approach for integrated testing and debugging and, in Sec. 4 we illustrate the approach through a working example. Then, in Sec. 5, we describe how Deipa was implemented as a Fiddle client tool, and we present conclusions and future work in Sec. 6.

2 The Fiddle Distributed Debugging Engine

Fiddle [LC00, LC01] is a distributed debugging engine based on a client-server operational model. The debugging engine acts as a server and the debugging tools, which are not included in Fiddle specification, act as the client(s).

The Fiddle debugging engine has the following main characteristics:

- *Support for interactive correctness debugging.* Fiddle includes all the traditional debugging facilities available in sequential debuggers, such as DBX [Lin90] and GDB [Sta88], but extended to operate upon the multiple processes of a distributed program. Some of the services provided by Fiddle operate at process level, e.g., breakpointing and single-stepping, while some other operate at application level, e.g., obtaining control of newly created application processes (*spawn'ed* in a PVM program). As a whole, Fiddle provides a consistent and complete set of services which would be unavailable if a simple set of independent sequential debuggers would have been used instead;
- *Client/server model.* Multiple debugging interfaces (client tools) can connect simultaneously to Fiddle and have access to the same set of target processes, providing, in this way, multiple “debugging views” over the same target application;
- *Extensibility.* Fiddle functionality may easily be extended by adding new clients, which provide specific and/or complementary functionalities to the debugging system;
- *Support of high-level user-defined abstractions.* Besides the ability to debug distributed programs at textual (source) level, Fiddle can also be applied to programs developed using higher level programming languages and/or models, such as visual parallel programming languages, which allow a parallel application to be specified in terms of a set of graphical entities describing application components and their interconnections. Fiddle extensions can be incorporated into the debugging system, which maps the functionalities required by such high-level debugging interface onto Fiddle basic debugging services;
- *Tool synchronization.* By allowing multiple concurrent client tools to access a common set of target processes, Fiddle needs to provide some basic support for tool synchronization. Using

Fiddle this can be achieved in two ways: i) as all requests made by client tools are controlled by a central daemon, Fiddle is able to serialize the service requests, avoiding some of the interferences among these tools; ii) client tools can receive notification events reporting changes in the target processes and/or service requests issued by other clients, allowing a tool to be aware of the “debugging environment”, and to react upon changes².

- *Easy integration in Parallel Software Engineering Environments.* Fiddle extensibility can be explored to support the cooperation and integration of debugging and other related tools in the environment, e.g., on-line program visualization tools.

2.1 Fiddle Software Architecture

Fiddle is structured as a hierarchy of five functional layers, each providing a specific set of debugging services and accessible through an interface library. Any layer may be used directly by a client tool if its set of services is found to be adequate for the client needs. Each layer is also used indirectly by the layer immediately above. For example, in Fig. 1, Layer 2_m has two clients: the tool $CT_1^{2_m}$ and Layer 3_m .

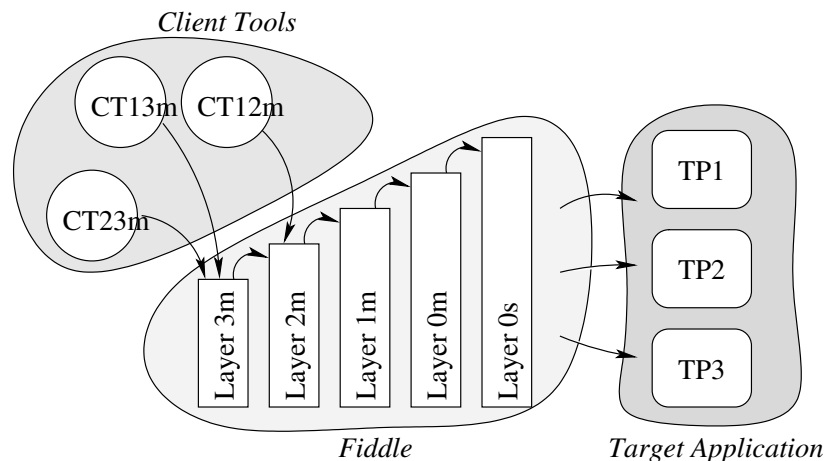


Figure 1: Fiddle’s Layered Architecture

A service requested by, e.g., a Layer 3_m client tool, will be processed and transferred successively to the underlying layers, until Layer 0_s is reached, and the service requested applied to the target process. The reply to such service request is also successively processed and transferred to the upper layers until the client tool gets the result.

There is a minimum set of functionalities common to all Fiddle layers, namely:

- *Inspect/control multi-threaded target processes.* Fiddle provides a basic set of debugging services to act upon threads within a multi-threaded process;
- *Inspect/control multiple target processes concurrently.* Any client tool may use Fiddle services to inspect and control multiple processes concurrently;
- *Support for client tool(s).* All layers accept, at least, requests from one client tool or from the layer above. The upper layers also accept multiple client tools operating concurrently upon the same target application.

Besides these common functionalities, each layer provides a set of specific functionalities. These functionalities are supported by the software architecture shown in Fig. 2.

²This coordination mechanism is not implemented in Fiddle yet. See Sec. 2.2.

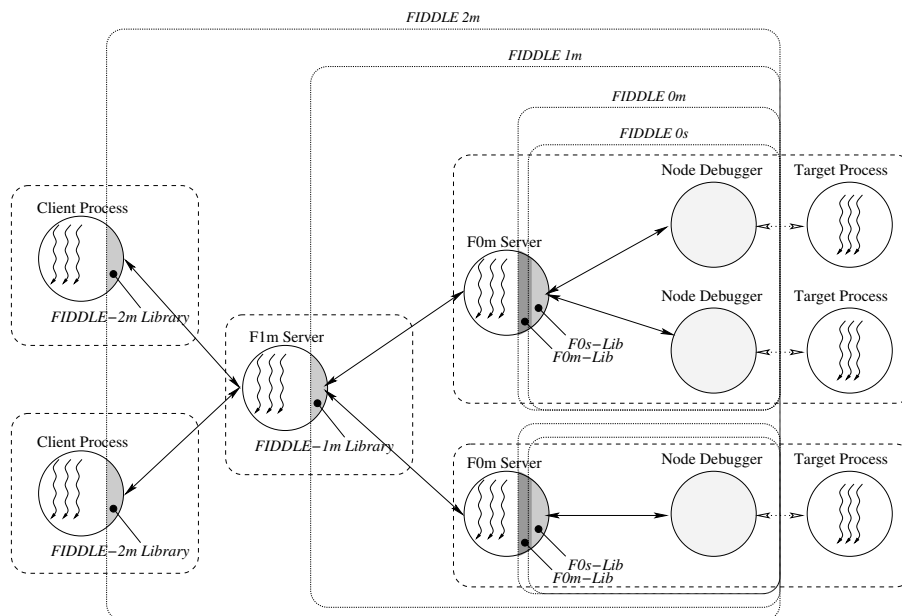


Figure 2: The Fiddle software architecture

Fiddle layers numbers range from *zero* to *three*, with incremental functionality, and have a suffix *s* or *m*, representing the class of accepted clients: *single-* or *multi-threaded*, respectively.

Layer 0_s This layer implements a set of local debugging services, including the common set of debugging services described above, which are made available to a single client tool.

The following components are known to this layer: i) *Target processes*: a subset of processes that compose the target application and are being executed in the local node; ii) *Node debuggers*: associated to each target process there is a node debugger, e.g., a GDB instance, which is responsible for acting upon it; iii) *Layer 0_s library*: provides single-threaded access to the debugging functionalities; iv) *Client tool*: a single-threaded tool accessing Fiddle and possibly implementing a user interface.

Layer 0_m This layer extends Layer 0_s to provide support for a single multi-threaded client tool. The client may now use threads to concurrently control multiple target processes and the interaction with the user.

Layer 1_m Both Layers $0_{s/m}$ have access to processes on the local node. Layer 1_m extends Layers $0_{s/m}$ to allow a client tool to have access to remote target processes.

The components known to this layer are: i) *Target processes*: running in the local or remote nodes; ii) *Node debuggers*: one associated to each target process; iii) *L 0_m Server*: an instance of this daemon will be automatically launched in each node hosting target processes; iv) *Layer 1_m Library*: provides thread-safe access to local and remote debugging functionalities; v) *Client tool*: client tools using this layer may transparently access local and remote target processes by using their global identifiers.

Layer 2_m This layer extends Layer 1_m to provide support for multiple simultaneous client tools, which may issue concurrent requests to Fiddle. This is achieved by using a *L 1_m Server* daemon, as an intermediary between the client tools and the *L 0_m Server* daemons. This *L 1_m Server* daemon multiplexes the service requests from the multiple client tools, submits them to Layer 1_m , and demultiplexes the corresponding replies back to the clients.

Layer 3_m Layers 0/1/2 services use a blocking semantics, keeping the calling thread waiting for the reply. This layer extends Layer 2_m to provide an event-based interaction between Fiddle and its clients. In contrast to the lower layers, a thread that invokes a method in the *Layer 3_m Library* doesn't block waiting for its reply. Instead, it receives a *Request Identifier* which will be used later to request and process the reply. When the service is executed, its success status and data (the *reply*) is sent to the client as an event. These events may be processed by the client tool in two different ways: i) *Asynchronous mode*: the general methodology to process Fiddle events is to define a handler. When a service is executed by Fiddle, a new thread will be created to execute the handler. A structure describing the results of the requested service is passed as an argument to the handling function, together with the *Request Identifier*; ii) *Synchronous mode*; in this mode, Fiddle keeps the notification of the event pending until the client tool explicitly requests it by invoking a Fiddle primitive.

2.2 Fiddle Current Status

Currently, Fiddle Layers 0/1/2_{s/m} are fully implemented and functional. Implementation of Layer 3_m, which will provide event notification and event handling mechanisms for client tools, is ongoing work. Fiddle development has taken part on Linux machines, but there are no strong dependencies on this operating system. Also, as the Fiddle debugging engine has no user interface, there are no dependencies on graphical environments/packages either.

If the target application is using PVM to spawn (launch) new processes, Fiddle can also automatically acquire such newly created processes, launching them under the control of a node debugger, and making them possible targets for future services requests.

Along with the Fiddle debugging engine, a set of text oriented debugging interfaces (Fiddle consoles) are being developed, one for each Fiddle layer, which allow to explore the full set of services provided by each layer. Additionally, other Fiddle client tools have been developed, such as FGI [Aug02] (the Fiddle Graphical Interface) based in the Gnome Desktop Environment, PADI [SNC00] (a group oriented parallel debugger) based in the Java AWT and Deipa (described in this paper).

Being still in a prototype stage, Fiddle is currently not publicly available (yet). However, it is available on a demand basis to anyone who request it to the author³.

3 Integrating Testing and Debugging: the DEIPA Approach

Program correctness requires that a given specification of the intended behavior be satisfied. Although the development of high-level abstractions for distributed programming has contributed to ease the task of program development, there are still many opportunities for programming errors, posing the need for support tools.

Many programming errors can be detected, in a more or less automatic way, through a static analysis of the program source text. Such analysis can also assist the programmer in the prediction of the program behavior concerning specific correctness properties. However, the program source text is not always available. Also, the global program behavior is often the result of a combination of behaviors, depending on the operating system or runtime support layers of a computing environment.

This explains the significance of approaches for dynamic analysis, which are centered upon the observation of real execution. Such approaches assume, from the beginning, the incomplete nature of this process, due to the usually huge number of computation states which can be generated by distributed program execution. As such, they are typically based on the selection of a finite set of representative test configurations, followed by an observation of the results of program execution. The definition of such test scenarios is, of course, dependent upon the classes of errors or program properties that one is trying to check.

³E-mail: jml@di.fct.unl.pt

The main goal of a debugging tool is to help analysing erroneous program behavior, possibly identified by a previous testing stage, and to assist the programmer in the formulation or confirmation of hypotheses on the causes of errors, and in the tracing of their origins in the program text.

Testing and debugging are naturally intertwined. On one hand, testing helps identifying errors whose causes must later be traced in a debugging stage. On the other hand, after a successful debugging session, one typically needs to reconsider the set of test configurations. Debugging can also help identifying unforeseen situations which may require the design of new testing scenarios.

The above aspects have been recognized for a long time, leading to many proposals of methodologies and tools for combined testing and debugging. See [KW99, CKW00] for a more complete survey on this topic.

Deipa [LCK⁺97] was initially developed to support a testing-and-debugging development cycle, by allowing the composition of two separately developed tools, STEPS [KW96] and DDBG [CLA99, CLD98]. The current configuration of such testing and debugging architecture is illustrated in Fig. 3.




Figure 3: Tool composition of STEPS and Fiddle using Deipa

Based on a static and dynamic analysis of the target program, STEPS generates a behavior specification file, the TeSS file. This file includes the definition of a list of *global breakpoints* (a consistent collection of local breakpoints). Some of those breakpoints have special instructions to modify one or more variable values. This is necessary to drive correctly the application execution when in presence of conditional branches or loops, e.g., on a `if` statement.

To force the application execution to conform to the specification in the TeSS file, Deipa generates sequences of debugging (control and inspection) commands to Fiddle, setting a local breakpoint in each target process and individually driving them until the breakpoints are reached. Once a process is stopped in a breakpoint, its internal status (e.g., variable contents) may be changed, according to the TeSS specification.

At any time, the Fiddle ability to handle multiple clients may be explored, and the user may switch to another tool, such as Fiddle's debugging console or graphical interface (FGI [Aug02]) and perform a more detailed inspection and control of the target processes. This allows the user to perform debugging at distinct abstraction levels, e.g., at a global program level to inspect process interactions, and at local process level, to inspect individual process execution.

On completing such a fine (process-level) debugging, the user may return to Deipa and proceed with the controlled execution mode, or release/stop the target application if no more debugging is needed.

4 A Working Example

To demonstrate how Deipa works, we have built a simple application. It is composed of two programs: the `echo_client`, which sends a number to the server and waits for its reply; and the

`echo_server`, which gets a number from the client and exits if the received number is even, or sends -1 back otherwise.

Since it is common for distributed applications to have a program that initializes the environment, we consider that the `echo_client` is the startup program, launching the `echo_server` server via `pvm_spawn()`.

First, we shall present the source code of both PVM programs, with `echo_client` on the left column and `echo_server` on the right one. The tiny leftmost number indicates the line number, used for future reference.

<code>echo_client</code>	<code>echo_server</code>
1 #include <pvm3.h>	1 #include <pvm3.h>
2 #include <stdio.h>	2 #include <stdio.h>
3 #include <stdlib.h>	3 #include <stdlib.h>
4 #include <unistd.h>	4
5	5
6	6 int main ()
7	7 {
8 int main ()	8 int mytid;
9 {	9 int from, value;
10 int mytid;	10
11 int totid=0;	11 mytid = pvm_mytid();
12 int value;	12
13	13 pvm_rcv (-1, -1);
14	14 pvm_upkint (&from, 1, 1);
15 mytid = pvm_mytid ();	15 pvm_upkint (&value, 1, 1);
16	16
17 value = pvm_spawn ("echo_server", NULL,	17 if ((value%2)==0)
18 PvmTaskDebug, ".", 1, &totid);	18 exit (0);
19	19 else
20 <i>/* Sending 0 will force the server to</i>	20 value=-1;
21 <i>exit, and the client will wait</i>	21
22 <i>forever for the reply */</i>	22 pvm_init send (0);
23	23 pvm_pkint (&value, 1, 1);
24 value = 0;	24 pvm_send (from, 1);
25 pvm_init send (0);	25
26 pvm_pkint (&mytid, 1, 1);	26 pvm_exit ();
27 pvm_pkint (&value, 1, 1);	27
28 pvm_send (totid, 1);	28 return 0;
29	29 }
30 <i>/* Get the reply from server */</i>	
31 pvm_rcv (-1, -1);	
32 pvm_upkint (&value, 1, 1);	
33 printf ("Received_value_%d\n", value);	
34	
35 pvm_exit ();	
36	
37 return (0);	
38 }	

The most relevant source lines in the client code are line 17, where the server is launched (spawned); line 28, where the message is sent to the server; and line 31, where the client blocks waiting for the reply.

On the server, the most relevant source lines are line 13, where the server waits for a message from the client; line 17, where the server decides how to react to the received message (terminate or reply); and line 24, where the server replies to the client (if this was the case).

As the development of the STEPS tools was discontinued by its authors, and Deipa uses TeSS files as one of its inputs, we produce the following TeSS files ourselves, which describe a testing scenario for the echo client and server processes. The tiny line numbers are not part of the TeSS file itself but were included only for future reference to specific lines.

```

1 START_FILE:
2     echo_client
3
4 SPAWN_TABLE:
5     {
6         0  0  0  0  1  echo_client echo_client.c 13142 4,
7         1 16  0  1  2  echo_server echo_server.c 59634 2
8     }
9
10 INITIAL:
11     [{ (1,1,17) }],
12     [{ (1,1,28) }],
13     [{ (2,1,28) }],
14     [{ (2,1,28),(1,2,13) }],
15     [{ (2,1,28),(2,2,13) }],
16     [{ (2,1,28),(1,2,17,[2,2,"value","1"]) }],
17     [{ (2,1,28),(2,2,17) }],
18     [{ (1,1,31),(1,2,24) }],
19     [{ (2,1,31),(2,2,24) }],
20     [{ (1,1,35),(1,2,26) }],
21     [{ (2,1,35),(2,2,26) }];

```

Each line represents a global breakpoint, as a collection of local breakpoints of the form “(T, I, L)”, which are the type of breakpoint, the virtual identifier of the process (VID) and the line number respectively. The type of breakpoint indicates if the program should stop before (code 1) or after (code 2) the selected line number. The virtual identifier of the target process will be mapped by Deipa to PVM *task IDs* at runtime. The new process, at line 14 of the TeSS file, is captured by the *launcher* program and integrated into the Fiddle and Deipa execution environment, as described in Sec. 5.

The TeSS file specifies global breakpoints in the following points:

- `Pvm_spawn()`, to ensure that new processes are, indeed, created on the PVM environment;
- Send and receive messages, to enable the verification of emission and reception of messages;
- `Pvm_exit()`, to ensure that all processes get out correctly.

4.1 Using Deipa and Fiddle to capture errors

The main idea of this example is to launch the distributed application and control its execution from the Deipa console; and to examine individual processes state on the Fiddle text user interface (f2m console). We must notice that this example requires the usage of *Layer 2_m*, as there are multiple client tools simultaneously connected to Fiddle, namely Deipa, the *launcher* and a Fiddle console.

The screen shots in Figures 4, 5 and 6 refer respectively to the following situations:

- The capture of the `echo_server` target process, created at line 17 of `echo_client`, and its integration into Deipa and Fiddle;
- Driving the application behavior by forcing its execution path; and
- The reception of the answer at the client side.

4.1.1 The capture of the new process

This task consists of intercepting the spawn action, the launching of the new target processes under the control of a node debugger, and the notification of Deipa and Fiddle on the existence of the new target process.

Figure 4 shows three windows, the PVM console on the top left, the Fiddle Layer 2_m console below, and Deipa on the right. On each window two states are shown. On the first state, the `echo_server` process hasn't been launched yet; while on the second, when the new process was already started (although not running).

```

[vrml@localhost tmp/aadebug/final_version/source_code/working_ ]$ pvm
pvm> ps
ps
  localhost.localdomain  40002  FLAG 0x COMMAND
  localhost.localdomain  40002  4/c -
pvm> ps
ps
  localhost.localdomain  40002  FLAG 0x COMMAND
  localhost.localdomain  40002  4/c -
  localhost.localdomain  40003  2/f echo_server
pvm> []

[vrml@localhost tmp/aadebug/final_version/source_code/working_ ]$ f2m
-----
| FIDDLE - v0.5.1      (Level 2m console) |
|-----|
| Copyright (c) 1999 DI - FCT/UNL      |
|-----|
| Type 'help' on prompt to get help    |
| <TAB> expands commands                |
| Abbreviations are accepted if unambiguous |
|-----|

f2m []> tids
TID  ATT  TP_PID  LLD_PID  L_TID MACHINE
  1    n    0       3108    1 localhost.localdomain

f2m []> tids
TID  ATT  TP_PID  LLD_PID  L_TID MACHINE
  1    n    0       3108    1 localhost.localdomain
  2    n    0       3138    2 localhost.localdomain

f2m []> []

[vrml@localhost working_example]$ ./console_deipa echo_example.tes
-----
| DEIPA - v0.1.0      (Level 2m console) |
|-----|
| Copyright (c) 2002 DI - FCT/UNL      |
|-----|
| Type 'help' on prompt to get help    |
| <TAB> expands commands                |
| Abbreviations are accepted if unambiguous |
|-----|

console_deipa > run
console_deipa > step
pth_manager: *** begin of process threads' list ***
( 1) prev=[bptype= 1 line= 17] actual=[bptype= 1 line= 28]
pth_manager: *** end of process thread's list ***
pth_launcher: connection from 127.0.0.1.32834, fd:7
pth_launcher: /home/vrml/pvm3/bin/LINUX/echo_server
pth_launcher: [tid=2, vid=2]

console_deipa > step
pth_manager: *** begin of process threads' list ***
( 1) prev=[bptype= 1 line= 28] actual=[bptype= 2 line= 28]
( 2) prev=[bptype=NULL line=NULL] actual=[bptype=NULL line=NULL]
pth_manager: *** end of process thread's list ***

console_deipa > []
  
```

Figure 4: The capture of a new process

Two of the commands available on the Deipa console are: `run`, to start execution the distributed application, and `step`, to advance the distributed application to the next global breakpoint indicated in the TeSS file.

On Deipa console, the distributed application (only `echo_client` for now) starts running and stops before line 17, the `pvm_spawn()` function call. When we step into the next global breakpoint, the `echo_client` process is spawned. As the `step` command lists all the known processes before doing any action, it's possible to verify that at the first `step` there is just one process (corresponding to the `echo_client` program), while at the second `step` there are already two processes, meaning the `echo_server` was already launched and is now under control of Fiddle (and Deipa).

Since Layer 2_m level allows several concurrent clients to connect to Fiddle, we also launched a Fiddle console. Once again, in the initial state (before the `pvm_spawn()`), there is only one process (with TID 1), but in the second state there are two processes (with TIDs 1 and 2), corresponding to `echo_client` and `echo_server`, respectively.

Similarly, in the top left window (the PVM console), one can notice that in the initial state there

is a single PVM task, and after the second step there are two PVM tasks.

From the observation of the target application with these three different tools, we may confirm that the new process was launched correctly and integrated into the debugging environment of Fiddle and Deipa.

4.1.2 Driving the application behavior

As stated before, the server program will exit if it receives an even number, leaving the client blocked forever waiting for a non-existing reply. We can change this behavior and force the server to always reply to the client, no matter it receives an odd or an even, by forcing the contents of the *value* variable at line 17 of *echo_server* to always be an odd number.

The Fig. 5 shows the same three windows as before. The PVM console on the left, the Fiddle Layer 2_m console just below, and Deipa on the right.

```

[vrn@localhost/tmp/aadebug/final_version/source_code/working_...]$ pvm
pvm> ps
ps
  localhost.localdomain  40002  FLAG 0x COMMAND
                        4/c -
pvm> ps
ps
  localhost.localdomain  40002  FLAG 0x COMMAND
  localhost.localdomain  40003  2/f echo_server
pvm> []

[vrn@localhost/tmp/aadebug/final_version/source_code/working_...]$ f2m
-----
| FIDDLE - v0.5.1      (Level 2m console) |
|-----|
| Copyright (c) 1999 DI - FCT/UNL      |
|-----|
| Type 'help' on prompt to get help    |
| <TAB> expands commands               |
| Abbreviations are accepted if unambiguous |
|-----|

f2m []> tids
TID  ATT  TP_PID  LLD_PID  L_TID MACHINE
  1    n    0       2736     1 localhost.localdomain

f2m []> tids
TID  ATT  TP_PID  LLD_PID  L_TID MACHINE
  1    n    0       2736     1 localhost.localdomain
  2    n    0       2766     2 localhost.localdomain

f2m []> 2 evaluate value
=> $1 = 134516282 ← old value (not initialized)

f2m []> 2 evaluate value
=> $2 = 1 ← new value

f2m []> []

[vrn@localhost/tmp/aadebug/final_version/source_code/working_...]$ ./console_deipa echo_example.tes
-----
| DEIPA - v0.1.0      (Level 2m console) |
|-----|
| Copyright (c) 2002 DI - FCT/UNL      |
|-----|
| Type 'help' on prompt to get help    |
| <TAB> expands commands               |
| Abbreviations are accepted if unambiguous |
|-----|

console_deipa > run

console_deipa > step
pth_manager: *** begin of process threads' list ***
(1) prev=[bptype= 1 line= 17] actual=[bptype= 1 line= 28]
    pth_manager: *** end of process thread's list ***
    pth_launcher: connection from 127.0.0.1:32826, fd:7
    pth_launcher: /home/vrn/pvm3/bin/LINUX/echo_server
    pth_launcher: [tid=2, vid=2]

console_deipa > step
pth_manager: *** begin of process threads' list ***
(1) prev=[bptype= 1 line= 28] actual=[bptype= 2 line= 28]
(2) prev=[bptype=NULL line=NULL] actual=[bptype=NULL line=NULL]
    pth_manager: *** end of process thread's list ***

console_deipa > step
pth_manager: *** begin of process threads' list ***
(1) prev=[bptype= 2 line= 28] actual=[bptype= 2 line= 28]
(2) prev=[bptype=NULL line=NULL] actual=[bptype= 1 line= 13]
    pth_manager: *** end of process thread's list ***

console_deipa > step
pth_manager: *** begin of process threads' list ***
(1) prev=[bptype= 2 line= 28] actual=[bptype= 2 line= 28]
(2) prev=[bptype= 1 line= 13] actual=[bptype= 2 line= 13]
    pth_manager: *** end of process thread's list ***

console_deipa > step
pth_manager: *** begin of process threads' list ***
(1) prev=[bptype= 2 line= 28] actual=[bptype= 2 line= 28]
(2) prev=[bptype= 2 line= 13] actual=[bptype= 1 line= 17]
    pth_manager: *** end of process thread's list ***
    set_all_vars_func: setvar value=1 ← setting a new value for the
                                         "value" variable

console_deipa > []

```

Figure 5: Driving the application behavior

Deipa was used to control the application until global breakpoint at line 15 is reached. At this point, we may use Fiddle console to inspect the *value* variable, which is reported to have the garbage value of 134516282.

Lets step to the next (line 16 of TeSS file) global breakpoint using Deipa, and re-inspect the *value* variable. Although this variable has been initialized to zero at line 15 of *echo_server*, it should contain now the value 1. This is achieved by the special global breakpoint of TeSS line 16, which also defines the new contents of the *value* variable. The contents of the *value* variable are checked again on the Fiddle console, and the value 1 confirmed.

4.1.3 Receiving the answer at the client side

In order to verify that the client always receives a reply message, we need to execute line 31 of the client's code. If the server has exited, the client will remain blocked forever at this line.

As shown in Fig. 6, the Fiddle console on the left window informs that the client is stopped at line 32, and Deipa on the right window shows that the client stopped after executing line 31 (it is type 2 global breakpoint).

```

[vrn@localhost tmp/aadebug/final_version/source_code/working_]# pvm
pvm> ps
ps
localhost.localdomain 40002 4/c -
pvm> ps
ps
localhost.localdomain 40002 4/c -
localhost.localdomain 40003 2/f echo_server
pvm> []

FIDDLE - v0.5.1 (Level 2m console)
Copyright (c) 1999 DI - FCT/UNL
Type 'help' on prompt to get help
<TAB> expands commands
Abbreviations are accepted if unambiguous

f2m []> tids
TID ATT TP_PID LLD_PID L_TID MACHINE
1 n 0 2736 1 localhost.localdomain

f2m []> tids
TID ATT TP_PID LLD_PID L_TID MACHINE
1 n 0 2736 1 localhost.localdomain
2 n 0 2766 2 localhost.localdomain

f2m []> 2 evaluate value
=> $1 = 134516282
f2m []> 2 evaluate value
=> $2 = 1
client received the message
f2m []> 1 info-line
Line 32 of "echo_client.c" starts at address 0x8049441 <main+233>
and ends at 0x8049454 <main+252>.
<</tmp/aadebug/final_version/source_code/working_example/echo_client.c:32:537:be
8049441
f2m []> []

console_deipa > step
pth_manager: *** begin of process threads' list ***
( 1) prev=[bptype= 1 line= 17] actual=[bptype= 1 line= 28]
pth_manager: *** end of process thread's list ***
pth_launcher: connection from 127.0.0.1,32826, fd:7
pth_launcher: /home/vrn/pvm3/bin/LINUX/echo_server
pth_launcher: [tid=2, vid=2]

console_deipa > step
pth_manager: *** begin of process threads' list ***
( 1) prev=[bptype= 1 line= 28] actual=[bptype= 2 line= 28]
( 2) prev=[bptype=NULL line=NULL] actual=[bptype=NULL line=NULL]
pth_manager: *** end of process thread's list ***

console_deipa > step
pth_manager: *** begin of process threads' list ***
( 1) prev=[bptype= 2 line= 28] actual=[bptype= 2 line= 28]
( 2) prev=[bptype=NULL line=NULL] actual=[bptype=NULL line=NULL]
pth_manager: *** end of process thread's list ***

console_deipa > step
pth_manager: *** begin of process threads' list ***
( 1) prev=[bptype= 2 line= 28] actual=[bptype= 2 line= 28]
( 2) prev=[bptype= 1 line= 13] actual=[bptype= 2 line= 13]
pth_manager: *** end of process thread's list ***

console_deipa > step
pth_manager: *** begin of process threads' list ***
( 1) prev=[bptype= 2 line= 28] actual=[bptype= 2 line= 28]
( 2) prev=[bptype= 1 line= 17] actual=[bptype= 2 line= 17]
pth_manager: *** end of process thread's list ***
set_all_vars_func: setvar value=1

console_deipa > step
pth_manager: *** begin of process threads' list ***
( 1) prev=[bptype= 2 line= 28] actual=[bptype= 2 line= 28]
( 2) prev=[bptype= 1 line= 17] actual=[bptype= 2 line= 17]
pth_manager: *** end of process thread's list ***

console_deipa > step
pth_manager: *** begin of process threads' list ***
( 1) prev=[bptype= 2 line= 28] actual=[bptype= 1 line= 31]
( 2) prev=[bptype= 2 line= 17] actual=[bptype= 1 line= 24]
pth_manager: *** end of process thread's list ***
client receives the message

console_deipa > step
pth_manager: *** begin of process threads' list ***
( 1) prev=[bptype= 1 line= 31] actual=[bptype= 2 line= 31]
( 2) prev=[bptype= 1 line= 24] actual=[bptype= 2 line= 24]
pth_manager: *** end of process thread's list ***
server sends the message

console_deipa > []

```

Figure 6: Receiving the answer at the client side

The above example shows how a set of separately developed tools: Deipa and Fiddle (and also the PVM console) can be used to drive a distributed application according to a specification of its desired behavior, which may differ from the application behavior in an uncontrolled environment.

5 Deipa as a Fiddle Client

Above Layer 2_m, Fiddle's API supplies functions to manage multiple clients on a single Fiddle session, interacting concurrently with Fiddle (and with the target application). Deipa, Fiddle's Layer 2_m console and FGI are examples of such clients.

5.1 The DEIPA tool

The Deipa tool is composed of two programs: i) *the main program*, which interacts with the user and controls the application behavior; and ii) *the launcher*, which captures newly created application processes (currently supporting only PVM) and makes their registration on Deipa and Fiddle.

The main program From Fiddle perspective, Deipa is just another client tool accessing the target application. To become a Fiddle client, and able to access its services, Deipa must be linked to the Fiddle Layer 2_m library and call the registration function available in the library API.

The launcher To capture new PVM processes and integrate them in the debugging environment (Deipa and Fiddle), the `PvmTaskDebug` flag must be activated when the target application spawns a new process, and the `PVM_DEBUGGER` environment variable must specify the *launcher* location. In this way, instead of spawning the new process directly, PVM spawns the *launcher*, which receives the name of the process to launch as a command line argument (for more details on this topic, consult the PVM documentation [GBD⁺94]).

Once the *launcher* is running, it executes the following steps:

1. Creates a bi-directional communication channel, for the future interactions with Fiddle;
2. Announces itself to Fiddle as a node debugger, indicating that future interactions should use the communication channel;
3. Announces itself to Deipa, informing on the new target process name and Fiddle identification number;
4. Redirects its standard I/O channels to the above mentioned communication channel;
5. Core image auto-replacement (`exec`) with a node debugger, running the new target process. The redirected I/O channels are inherited on `exec`'ing, and the node debugger will have its I/O redirected to the communication channel.

The Fiddle API includes a function that allows to register/incorporate an existing node debugger into the Fiddle environment. The Deipa *launcher* fits into this category, but before becoming a node debugger, it supplies some relevant information to Deipa.

5.2 Inside Deipa

Deipa is structured in two main parts: i) *the front-end*, currently only a text-oriented user interface (Deipa console) is implemented; and ii) *the back-end*, where all the relevant Deipa functionalities are implemented.

The Deipa console The commands available to the user can be divided in two sets: i) the ones whose scope is limited to Deipa itself, e.g., `open` (load a new TeSS file), and ii) those that act upon the target processes, e.g., `run` and `step` (drive the target application until the first/next global breakpoint).

The Deipa library Deipa library, whose architecture is described in Fig. 7, uses multiple threads, one for each target process (the *process threads*) and three internal management threads.

The process threads are responsible for the control and inspection of target processes.

The three management threads are: i) the *main thread*, responsible for receiving requests, verifying if a request is valid in the current state of the library, and sending the replies; ii) the *manager thread*, responsible for managing all the process threads; and iii) the *launcher thread*, responsible for receiving the registration of each new process created by the `pvm_spawn()` and signaled by the launcher program.

Details on Deipa internals can be found in [Mor02].

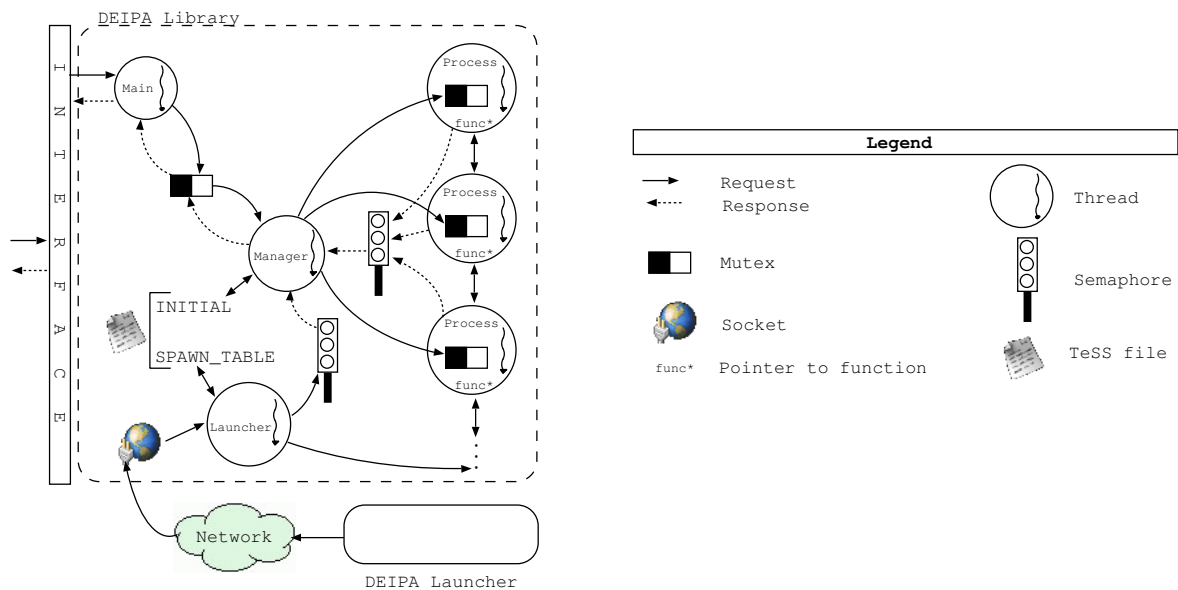


Figure 7: Deipa implementation

6 Conclusions and Future Work

The Fiddle engine meets several requirements for the debugging of multi-thread and multi-process distributed applications. Its suitability has been evaluated through the development of prototypes to test the integration of several types of support tools. This includes the support for graphic debugging interfaces, the integration of testing and debugging, and the integration of visualization and debugging tools.

In this paper, we have described how Fiddle eases the development of integrated testing and debugging support. This is achieved through a separate tool, Deipa, which acts as an intermediary between an independent testing tool (STEPS) and the Fiddle debugging engine. Furthermore, Fiddle provides support for concurrent clients, allowing the user to perform debugging and control of a distributed application at two distinct levels: the Deipa level, where global program states are followed, and at the Fiddle level, where local process states can be inspected and modified.

Ongoing work on Fiddle client tools, such as FGI and PADI, will lead to new experiments on tool integration and cooperation, with the increased user friendliness given by the graphical user interfaces. There are plans to do a similar work on Deipa, replacing its text console with a much friendlier graphical interface.

Communication between the client tools and Fiddle, and internally between Fiddle components, has been reformulated and changed from a proprietary data format to XML. This usage of the XML standard will improve Fiddle ability to integrate with third party software packages.

References

- [Aug02] P. Augusto. Fgi — fiddle graphical interface. Technical report, Departamento de Informática, Universidade Nova de Lisboa, December 2002.
- [CD98] J. C. Cunha and V. Duarte. Monitoring PVM programs using the DAMS approach. *Lecture Notes in Computer Science*, 1497:273–280, 5th Euro PVM/MPI 1998.

- [CKW00] J. C. Cunha, P. Kacsuk, and S. Winter, editors. *Parallel Program Development for Cluster Computing: Methodology, Tools and Integrated Environment*. Nova Science Publishers, Inc., 2000.
- [CLA99] J. C. Cunha, J. Lourenço, and T. Antão. An experiment in tool integration: the DDBG parallel and distributed debugger. *Euromicro Journal of Systems Architecture*, 45(11):897–907, 1999. Elsevier Science Press.
- [CLD98] J. C. Cunha, J. Lourenço, and V. Duarte. Using DDBG to support testing and high-level debugging interfaces. *Computers and Artificial Intelligence*, 17(5):429–439, 1998. Slovak Academic Press.
- [CLV⁺98] J. C. Cunha, J. Lourenço, J. Vieira, B. Moscão, and D. Pereira. A framework to support parallel and distributed debugging. In *Proceedings of the International Conference on High-Performance Computing and Networking (HPCN'98)*, Amsterdam, The Netherlands, April 1998.
- [GBD⁺94] A. Geist, A. Beguelin, J. Dongarra, W. Jiang, R. Manchek, and V. S. Sunderam. *PVM: Parallel Virtual Machine: A Users' Guide and Tutorial for Networked Parallel Computing*. Scientific and Engineering Computation. MIT Press, Cambridge, MA, USA, 1994.
- [Hoo96] Robert Hood. The p2d2 project: Building a portable distributed debugger. In *Proceedings of the 2nd Symposium on Parallel and Distributed Tools (SPDT'96)*, Philadelphia PA, USA, 1996. ACM.
- [KGV96] D. Kranzlmuller, S. Grabner, and J. Volkert. Debugging massively parallel programs with ATTEMPT. In H. Liddell, A. Colbrook, B. Hertzberge, and P. Sloot, editors, *High-Performance Computing and Networking (HPCN'96 Europe)*, volume 1067 of LNCS, pages 798–804. Springer-Verlag, 1996.
- [KP00] J. M. Kewley and R. Prodan. A distributed object-oriented framework for tool development. In *Proceedings of the 34th International Conference on Technology of Object-Oriented Languages and Systems (TOOLS 34)*, pages 353–362. IEEE Computer Society Press, July 2000.
- [KW96] H. Krawczyk and B. Wiszniewski. Interactive testing tool for parallel programs. In I. Jelly, I. Gorton, and P. Crolll, editors, *Software Engineering for Parallel and Distributed Systems*, pages 98–109, London, UK, 1996. Chapman & Hal.
- [KW99] H. Krawczyk and B. Wiszniewski. *Analysis and Testing of Distributed Software Applications*. John Wiley & Sons, November 1999. ISBN 0471978027.
- [LC00] J. Lourenço and J. C. Cunha. *Flexible Interface for Distributed Debugging (Library and Engine): Reference Manual (V 0.3.1)*. Departamento de Informática da Universidade Nova de Lisboa, Portugal, December 2000. Under development.
- [LC01] João Lourenço and José C. Cunha. Fiddle: A flexible distributed debugging architecture. In *Proc. of the International Conference on Computational Science (ICCS'2001), Special Session on "Tools and Environments for Parallel and Distributed Programming"*, volume 2074, pages 821–830, San Francisco, CA, USA, May 2001. Springer-Verlag. ISBN 3-540-42233-1.
- [LCK⁺97] J. Lourenço, J. C. Cunha, H. Krawczyk, P. Kuzora, M. Neyman, and B. Wiszniewsk. An integrated testing and debugging environment for parallel and distributed programs. In *Proceedings of the 23rd Euromicro Conference (EUROMICRO'97)*, pages 291–298, Budapest, Hungary, September 1997. IEEE Computer Society Press.
- [Lin90] Mark A. Linton. The evolution of Dbx. pages 211–220, Summer 1990.

- [LWO96] T. Ludwig, R. Wismueller, and M. Oberhuber. OCM — an OMIS compliant monitoring system. *Lecture Notes in Computer Science*, 1156, 1996.
- [LWSB97] T. Ludwing, R. Wismüller, V. Sunderam, and A. Bode. OMIS – On-line monitoring interface specification. Technical report, LRR-Technish Universiät München and MCS-Emory University, 1997.
- [Mor02] V. Moreira. Deipa e codificação xml para o formato xdr do fiddle. Technical report, Departamento de Informática, Universidade Nova de Lisboa, July 2002.
- [SNC00] D. Stringhini, P. Navaux, and J. Chassin de Kergommeaux. A selection mechanism to group processes in a parallel debugger. In *PDPTA'2000*, Monte Carlo Resort, Las Vegas, Nevada, USA, June 2000.
- [Sta88] Richard Stallman. *GDB manual: the GNU source-level debugger*. Free Software Foundation, Inc., Cambridge? Mass., 2nd, GDB version 2.5 edition, February 1988.
- [Too] Dolphin ToolWorks. *TotalView*. Dolphin Interconnect Solutions, Inc., Framingham, Massachusetts, USA.
- [TSLW00] J. Trinitis, V. Sunderam, T. Ludwig, and R. Wismüller. Interoperability support in distributed on-line monitoring systems. In *Proc. of the International Conference on High-Performance Computing and Networking (HPCN'2000)*, volume 1823, Amsterdam, The Netherlands, 2000.

