

YAMI Core Library Tutorial

Copyright © 2001-2008 Maciej Sobczak

Contents

1	Introduction	2
2	Echo	2
2.1	Easy client	2
2.2	Easy client with data	3
2.3	Full blown client	4
2.4	Server, polling version	6
2.5	Server, passive version	8
2.6	Asynchronous vs. synchronous	10
2.7	Error handling	11
3	Calculator	11
3.1	Client	12
3.2	Server	14
3.3	Explicit event processing	17

1 Introduction

In this tutorial, there are two assumptions:

- Every server runs on the port 12340. This is arbitrary choice.
- Every client runs on the same machine as the server (this is for simplicity only – if you have access to the network, try the examples running the processes on different machines), and if it requires its own listening socket, its port is assigned by the operating system.

2 Echo

As a first example, you will write your own *echo* server and client. The echo server will wait for messages from clients and will print them on the screen. Apart from that, the server will also accept a special *shutdown* message. There will be no data sent back to clients.

2.1 Easy client

First, you will learn how to write simple client. In most distributed systems it is easier to write clients than servers.

The first client in this tutorial sends a message to the server that tells it to shut down. Here it is (file `shutdown.c`):

```
1 #include "yamic.h"
2 #include <stdlib.h>
3
4 int main()
5 {
6     yamiNetInitialize();
7
8     yamiEasySend("127.0.0.1", 12340, 2,
9         "echo", "shutdown", NULL);
10
11     yamiNetCleanup();
12
13     return 0;
14 }
```

Yes, it is *that* easy.

The first line is necessary to use *YAMI* in C.

The line with call to `yamiEasySend` alone is enough to send a message to the server. There are six parameters used in this function:

- "127.0.0.1" – the server's address. It is a string value and can have the form "comp.company.com" as well.

- 12340 – the server’s port
- 2 – the level of the message sent. In this example it can be 1 as well (which means “only strings, please”). In most cases you want it to be Level2, unless you are sending a message to the component that understands only strings.
- "echo" – the name of the destination object. In *YAMI*, messages are sent to objects which means that there can be many destinations in the same server process.
- "shutdown" – the message name. Later you will see that this message causes the server to shut down.
- NULL – not used now. In the next client example, you will use this parameter to pass some data along with the message.

Note also the `yamiNetInitialize` and `yamiNetCleanup` functions. They are needed on MS Windows system. If you do *not* plan to write for Windows, you can safely omit them (they are empty for Unix/Linux) but keep in mind that using them will add to portability of your code.

Interesting? So let’s pass some data with the message.

2.2 Easy client with data

In the previous section, the message had only a name. In many cases it is enough, because it allows the server to make decisions about what action to perform, but there are situations where it would be nice to pass some data together with the message. In *YAMI*, there are two levels of messages with regard to the data that can be sent:

- Level1 – this level allows only strings and wide strings. It was designed for the components written in languages that do not understand other data types.
- Level2 – this level allows strings, wide strings, integers, doubles, bytes and arbitrary binary blocks.

Both levels are accessible in *YAMI* Core Library.

There can be many values of different types that can go with the same message at once. All the values are stored in the so-called *Parameter Set*, which is basically a list of values of chosen type. Such a *Parameter Set* has to be created before the message is sent and appended to it. It can be even used many times with different messages, but we will not use this feature here.

The client that causes the server to print “Hello, YAMI!” message on the screen can look like here (file `hello.c`):

```
1 #include "yamic.h"
2 #include "yamiparams.h"
```

```

3
4 int main()
5 {
6     HPARAMSET params;
7
8     yamiNetInitialize();
9
10    yamiCreateParamSet(&params, 1);
11    yamiSetString(params, 0, "Hello, YAMI!");
12
13    yamiEasySend("127.0.0.1", 12340, 2,
14                "echo", "print", params);
15
16    yamiDestroyParamSet(params);
17
18    yamiNetCleanup();
19
20    return 0;
21 }

```

Yes, it is *that* easy.

The second line is necessary to allow you to use *Parameter Set* functionality. The line 6 defines a local variable that is a handle to the *Parameter Set* object. Later, the *Parameter Set* is created to contain 1 uninitialized parameter and after that, the parameter is set to be a string of value “Hello, YAMI!”. As you can see in the next line, the *Parameter Set* object is used to send the message. Note also that the message name has changed. At the end, the *Parameter Set* object is destroyed, since it will be no longer needed.

2.3 Full blown client

The examples written so far have one thing in common: they send only one message to the server and quit. There is nothing that prevents you from sending millions of messages using the `yamiEasySend` function, but in terms of performance it will not be a good idea. The `yamiEasySend` function creates a *YAMI Agent* that is responsible for actually sending a message and destroys it after that. If you want to send many messages, it is a better idea to follow the scheme:

1. Create the *Agent* object.
2. Send the message. Repeat this step as many times as you wish.
3. Destroy the *Agent*.

In this section you will write the client according to this scheme. The client itself will be smarter than before, too – it will read the lines of text from its standard input and send each line as a separate message to the *echo* server. The code is below (file `printall.c`):

```
1 #include "yamic.h"
2 #include "yamiparams.h"
3 #include <stdio.h>
4 #include <string.h>
5
6 #define MAXBUF 1000
7
8 void readline(char *buf)
9 {
10     int i;
11
12     fgets(buf, MAXBUF, stdin);
13     for (i = strlen(buf); i >= 0; --i)
14         if (buf[i] == '\n' || buf[i] == '\r')
15             buf[i] = 0;
16 }
17
18 int main()
19 {
20     HYAMIAGENT agent;
21     HPARAMSET params;
22     char buf[MAXBUF];
23
24     yamiNetInitialize();
25
26     yamiCreateAgent(&agent, 0, NULL);
27     yamiAgentDomainRegister(agent,
28         "echoserver", "127.0.0.1", 12340, 2);
29
30     readline(buf);
31     while (!feof(stdin))
32     {
33         yamiCreateParamSet(&params, 1);
34         yamiSetString(params, 0, buf);
35
36         yamiAgentMsgSend(agent, "echoserver",
37             "echo", "print", params, NULL);
38
39         yamiDestroyParamSet(params);
40
41         readline(buf);
42     }
43
44     yamiDestroyAgent(agent);
45     yamiNetCleanup();
46 }
```

```

47     return 0;
48 }

```

The HYAMIAGENT is a type of handle to the *Agent* object.

The *Agent* is created with its own listening port, which number is assigned by the operating system. This *Agent* does not use this port, since no information is sent back to the client. Please read about *Policies* to learn how to create the *Agent* object without the listening socket.

After the *Agent* is created, the remote domain is registered in it, like in *White Pages Book*. The address and port of destination *Agent* (the one created by the server) is remembered under the name "echoserver" for later use (the name in the address book is arbitrary).

The program performs a loop where it reads lines of text from standard input. Each line is then sent as a parameter with the message to the server. The last argument to the `yamiAgentMsgSend` function allows to retrieve responses from the server. This will be used in later examples.

2.4 Server, polling version

Finally, we will write the server.

Err... what does this *polling* mean? First, write the following code (file `servera.c`):

```

1 #include "yamic.h"
2 #include "yamiparams.h"
3 #include <stdio.h>
4 #include <string.h>
5
6 int main()
7 {
8     HYAMIAGENT agent;
9     HINCMSG incoming;
10    HPARAMSET params;
11    const char *msgname;
12    const char *line;
13
14    yamiNetInitialize();
15
16    yamiCreateAgent(&agent, 12340, NULL);
17    yamiAgentObjectRegister(agent, "echo",
18        polling, NULL, NULL);
19
20    puts("server started");
21    while (1)
22    {
23        yamiAgentIncomingMsgGetNext(agent, "echo",

```

```

24         &incoming, 1);
25
26     yamiAgentIncomingMsgGetMsgName(incoming,
27         &msgname);
28
29     if (strcmp(msgname, "shutdown") == 0)
30     {
31         puts("received the shutdown message");
32         yamiAgentDestroyIncomingMsg(incoming);
33         break;
34     }
35     else /* if (strcmp(msgname, "print") == 0) */
36     {
37         yamiAgentIncomingMsgGetParameters(incoming,
38             &params);
39
40         yamiGetStringBuffer(params, 0, &line);
41         puts(line);
42
43         yamiDestroyParamSet(params);
44         yamiAgentDestroyIncomingMsg(incoming);
45     }
46 }
47
48     yamiDestroyAgent(agent);
49     yamiNetCleanup();
50
51     return 0;
52 }

```

The `HINCMMSG` is a type of handle to the incoming message.

After the *Agent* is created, the *polling* object is registered in it as a potential target of messages. If any message comes with the object name that is equal to the name of the registered object, the *Agent* will store the message in a queue, one queue for one object. Later, the server program goes into a loop that... asks the *Agent* if there are any messages for the `echo` object and processes them. This is why the server is *polling* – it has to ask the *Agent* for every new message. Moreover, in this example, the server will block waiting for new message if the queue is empty – the last parameter to `yamiAgentIncomingMsgGetNext` is non-zero.

When the message is finally retrieved, the server checks its name. If it is `"shutdown"`, it breaks the loop and terminates. Otherwise (which means that the message name is `"print"`; in real system you would also take some steps if the message has unknown or unacceptable name), the server gets the *Parameter Set* from the message and prints its first parameter on the standard output.

Note that both the *Parameter Set* and the incoming message handle are de-

stroyed when not needed.

2.5 Server, passive version

If there was a *polling* server, then there must be some alternative. It is called a *passive* server. Previously, the server had to ask for each incoming message. The advantage was that the server could ask in the most appropriate time if it had some other things to do. However, in some situations it is convenient if it is the *Agent* who tells the server that there is new message. In this scheme, the server itself is passive – it only waits for invocations from the *Agent*. It is important to note that the invocations are made from the separate thread that belongs to the *Agent* object. Let's see the code (file `serverb.c`):

```

1 #include "yamib.h"
2 #include "yamiparams.h"
3 #include "yamisynchro.h"
4 #include <stdio.h>
5 #include <string.h>
6
7 HSEMAPHORE sem;
8
9 void servant(HINCMSG incoming)
10 {
11     HPARAMSET params;
12     const char *msgname;
13     const char *line;
14
15     yamiAgentIncomingMsgGetMsgName(incoming, &msgname);
16     if (strcmp(msgname, "shutdown") == 0)
17     {
18         puts("received the shutdown message");
19         yamiSemaphoreRelease(sem);
20     }
21     else /* if (strcmp(msgname, "print") == 0) */
22     {
23         yamiAgentIncomingMsgGetParameters(incoming,
24             &params);
25
26         yamiGetStringBuffer(params, 0, &line);
27         puts(line);
28
29         yamiDestroyParamSet(params);
30     }
31
32     yamiAgentIncomingMsgEat(incoming);
33 }
```

```

34
35 int main()
36 {
37     HYAMIAGENT agent;
38
39     yamiNetInitialize();
40
41     yamiCreateSemaphore(&sem, 0);
42
43     yamiCreateAgent(&agent, 12340, NULL);
44     yamiAgentObjectRegister(agent, "echo",
45         passive_singlethreaded, servant, NULL);
46
47     puts("server started");
48
49     yamiSemaphoreAcquire(sem);
50
51     /* the process reaches this code      */
52     /* only when the semaphore is released */
53
54     yamiDestroyAgent(agent);
55     yamiDestroySemaphore(sem);
56     yamiNetCleanup();
57
58     return 0;
59 }

```

The `yamisynchro.h` file is included so that it will be possible to provide some synchronization between two threads involved here into processing.

The first difference in the `main` function is that the object was registered as a `passive_singlethreaded` instead of `polling`. The last but one parameter to the `yamiAgentObjectRegister` function is a pointer to function that will be called by the *Agent* when there is some new message for the given object. After registering the object, the server process suspends its execution with simple trick – acquire the global semaphore that is initially set to 0. In other words, the server process is going to sleep.

Now, let's go back to the `servant` function. The name was chosen because in other distributed infrastructures (notably CORBA) the servant is a piece of software that is responsible for processing messages. This is exactly the purpose of the function. Note that the `incoming` handle to the message is already provided. What does the `servant` function do? The same as before, with few differences:

- There is no loop to break – in order to wake up the main server thread, the global semaphore is released. This will cause the main thread to resume execution, clean up and terminate.

- The `incoming` handle is not destroyed – it is a property of the *Agent* object.
- The incoming message is marked as *eaten*. Without this, the *Agent* would try to send back to the client some information about the outcome of the processing and in this example the client does not want to retrieve anything. To be exact, the *Agent* expects that the message will be either processed or rejected, when “processed” means replied or eaten. Since there is no reply and no rejection, we mark the message as eaten.

2.6 Asynchronous vs. synchronous

YAMI sends and processes messages asynchronously by default. This means two things:

- When you “send” the message, it is only put into the sending queue. There is a separate sender thread in the *Agent* that is responsible for actually sending messages. The advantage? You can put many messages into the queue and go back to your work without waiting for them to be physically sent.
- When the message arrives at the remote *Agent*, it is put into the queue associated with each registered object and waits there for processing. It is natural with polling servers (the messages have to wait *somewhere* until you ask for them), but is also used by default for passive servers. The advantage? The message can be quickly received and stored while you are busy with long computations or processing previous messages.

The asynchronous sending and receiving is not always good, however. Let’s say that you run the command:

```
$ ./printall
```

(I assume some Unix environment here)

and later type some lines of text by hand. You are not really fast with typing, so the messages (each message contains one line of text) are actually sent to the server before you get your finger off the keyboard. But try this little experiment:

```
$ ./printall < printall.c
```

which should send as many messages to the server as there are lines in the `printall.c` file. After reaching end-of-file, the process terminates... which can happen before the sender thread sends anything! Or you will see only few of the lines on the server side. There are three solutions to this problem:

- Put some delay at the end of the client process, but before the *Agent* is destroyed, so that the sender thread will have a chance to send all messages from its queue. Do not tell anybody.

- Rewrite the client and server so that the server will send back some response after each message. Wait for this response before the client sends the next message. Do not tell anybody that the excessive network traffic is your fault.
- Create the *Agent* in the client process *without* the separate sender thread. In such configuration, every message will be sent immediately, without storing in any queue – it will be synchronous. This way you will insure that each line is physically sent to the remote end (to the server) before next line is read from the standard input. Show your code to everybody.

Similar issues arise on the server side. The messages received from the network are automatically stored in the queue. This queue can get overflowed (you can increase its maximum size, but again – do not tell anybody). Similarly, the professional solution is to change the dispatching mode to synchronous by switching off dispatching threads in the *Agent*.

Look at the `printsync.c` and `servsync.c` files to see how to do it. If you compile them, you can try the test:

```
$ ./servsync
```

in one console and (probably on different computer in the network, but you should change the address used in the client source code) in the other:

```
$ ./printsync < somelongfile
```

Enjoy.

2.7 Error handling

Many different things can go wrong in the distributed system. Some of them will be your fault, some not. Almost every *YAMI* function reports success or error condition. In the example programs presented so far there was no error checking – for simplicity and to avoid cluttering the source code. In real programs, you will for sure want to check for error codes. The `printsync.c` and `servsync.c` files are written with simple, but consistent error checking.

3 Calculator

In this second example, you will write the client-server pair where the client asks for something and the server *responds* to the messages, sending some data back to client. This is probably the most common way of developing distributed systems.

The calculator example consists of:

- The server that accepts messages: `add`, `sub`, `mul` and `div`, each with two integer parameters. Each message will be replied to, with the appropriate integer result.

- The client that sends messages asking the server to perform some computation.

I'm sure that after this example you will be able to write distributed systems of arbitrary complexity.

3.1 Client

Let's start with easy things. The client (file `calcclient.c`) looks like here:

```
1 #include "yamih.h"
2 #include "yamiparams.h"
3 #include <stdio.h>
4
5 const char *serverhost = "127.0.0.1";
6 const int serverport = 12340;
7 const int clientport = 0; /* use system assigned port */
8 const char *domainname = "someDomain";
9 const char *objectname = "calculator";
10
11 int main()
12 {
13     HYAMIAGENT agent;
14     HPARAMSET paramset, returnparamset;
15     HMESSAGE hm;
16     enum msgStatus status;
17     int result;
18
19     yamiNetInitialize();
20
21     yamiCreateAgent(&agent, clientport, NULL);
22
23     yamiAgentDomainRegister(agent, domainname,
24                             serverhost, serverport, 2);
25
26     yamiCreateParamSet(&paramset, 2);
27
28     yamiSetInt(paramset, 0, 100);
29     yamiSetInt(paramset, 1, 20);
30
31     yamiAgentMsgSend(agent, domainname,
32                     objectname, "add", paramset, &hm);
33
34     yamiDestroyParamSet(paramset);
35
36     yamiAgentMsgWait(hm);
37
```

```

38     yamiAgentMsgGetStatus(hm, &status);
39
40     if (status == eReplied)
41     {
42         yamiAgentMsgGetResponse(hm, &returnparamset);
43
44         yamiGetInt(returnparamset, 0, &result);
45
46         yamiDestroyParamSet(returnparamset);
47
48         printf("the result is %d\n", result);
49     }
50     else if (status == eRejected)
51     {
52         puts("the last message was rejected");
53     }
54     else
55     {
56         puts("no correct reply");
57     }
58
59     yamiAgentMsgDestroy(hm);
60     yamiDestroyAgent(agent);
61     yamiNetCleanup();
62
63     return 0;
64 }

```

The main difference to the previous example is the usage of the last parameter to the `yamiAgentMsgSend` and its consequences. The `HMESSAGE` is a type of handle to the message token. It allows you to retrieve some information about the message sent to the remote objects. It also allows you to *synchronize* the client activity with the server – the `yamiAgentMsgWait` function allows the client to wait for some change in the message's status.

As you can see, the message is sent with two integer parameters in the *Parameter Set* (100 and 20). The message is sent asynchronously, so that you can continue your job without waiting for response and ask for it later. In this example there is nothing interesting to do anyway, so we decide to *wait* (with the call to `yamiAgentMsgWait`) until something interesting happens to the message token. The process wakes up when (for example) the response arrives. The `status` of the message is examined to find out if there is a real reply or maybe some network error or something else. If there is a reply, we just retrieve the returning *Parameter Set* and print its only (first) parameter.

The Calculator example shows also that there are two contexts where *Parameter Set* can be used:

- When the client sends a message to the server. The *Parameter Set* can

be supplied at the client-side and retrieved at the server-side.

- When the server replies to the message. The server can supply a *Parameter Set* and the client can extract it after receiving the response.

These two contexts allow to send data in both directions.

Easy? So take also a look at the `calcclient2.c` file and play with it – it is an interactive calculator console.

There is one thing to remember, though. The `yamiAgentMsgWait` function is your friend, but do not trust him. It may happen that the server crashes in the middle of the computations. Then – the client does not receive any notification and the status of the message is always `ePending`. If you call the wait function, you are in troubles (the other solution, sometimes reasonable, is to periodically ask the message token if the response arrived – the client can decide by himself that something went wrong after, say, 10th try). To help you with this problem, the *Agent* object provides the *waker* service. The `calcclient2.c` file shows how to use it.

3.2 Server

The server is presented below (file `calcserver.c`):

```

1 #include "yamic.h"
2 #include "yamiparams.h"
3 #include "yamisynchro.h"
4 #include "yamierrors.h"
5 #include <stdio.h>
6 #include <string.h>
7
8 const int  serverport  = 12340;
9 const char *objectname = "calculator";
10
11 void servant(HINCMSG incomingmsg)
12 {
13     HPARAMSET paramset, returnparamset;
14     const char *msgname;
15     int val1, val2, result;
16
17     yamiAgentIncomingMsgGetMsgName(incomingmsg,
18                                     &msgname);
19
20     printf("I have received the message: %s\n",
21           msgname);
22
23     if (strcmp(msgname, "add") &&
24         strcmp(msgname, "sub") &&
25         strcmp(msgname, "mul") &&

```

```
26     strcmp(msgname, "div"))
27     {
28         puts("unknown name - rejecting");
29         yamiAgentIncomingMsgReject(incomingmsg);
30         return;
31     }
32
33     yamiAgentIncomingMsgGetParameters(incomingmsg,
34         &paramset);
35
36     yamiGetInt(paramset, 0, &val1);
37     yamiGetInt(paramset, 1, &val2);
38
39     yamiDestroyParamSet(paramset);
40
41     if (strcmp(msgname, "add") == 0)
42     {
43         result = val1 + val2;
44     }
45     else if (strcmp(msgname, "sub") == 0)
46     {
47         result = val1 - val2;
48     }
49     else if (strcmp(msgname, "mul") == 0)
50     {
51         result = val1 * val2;
52     }
53     else /* msgname is "div" */
54     {
55         if (val2 == 0)
56         {
57             puts("dividing by 0 not allowed");
58             puts("rejecting");
59             yamiAgentIncomingMsgReject(incomingmsg);
60             return;
61         }
62         else
63             result = val1 / val2;
64     }
65
66     yamiCreateParamSet(&returnparamset, 1);
67
68     yamiSetInt(returnparamset, 0, result);
69
70     yamiAgentIncomingMsgReply(incomingmsg,
71         returnparamset);
```

```

72
73     yamiDestroyParamSet(returnparamset);
74 }
75
76 int main()
77 {
78     HYAMIAGENT agent;
79
80     puts("starting the server");
81
82     yamiNetInitialize();
83
84     yamiCreateAgent(&agent, serverport, NULL);
85
86     yamiAgentObjectRegister(agent, objectname,
87         passive_singlethreaded, servant, NULL);
88
89     puts("going to sleep...");
90     yamiSleep(0);
91
92     yamiDestroyAgent(agent);
93
94     yamiNetCleanup();
95
96     return 0;
97 }

```

This is a *passive* server, because it uses the *call-back* pattern for message dispatching. You have seen it already in the Echo example.

This server has one design quirk. I show you this, because it is sometimes used in servers working in distributed systems (I have seen it in some CORBA examples, really). Namely – after setting up everything in the `main` function, the main server thread goes to sleep... and never wakes up. The `yamiSleep` function with 0 as a parameter sleeps *forever*. This means two things:

- There is no way to remotely shut the server down. Locally, the administrator can just kill the server process and bless the operating system that will hopefully garbage-collect all the resources used by the process.
- All the code after the sleep function is nothing but a bluff. It will never execute.

I do not advocate this style of writing servers, although, as I've pointed out, I *have* seen it. The two clean solutions that I recommend are:

- Provide additional shut-down message, like in the Echo example. Synchronize servant function with the main thread if the server is passive or just break the message loop if it is polling.

- Provide (register) additional shut-down *object* and after setting up everything in the `main` function, call the `yamiAgentIncomingMsgGetNext` function on this object instead of going to sleep. The first message that will be received for this special shut-down object will wake you up, so that you can terminate the server in a clean way.

The file `calcserver2.c` is the same as presented above `calcserver.c` but with error-handling added.

3.3 Explicit event processing

This part of the tutorial is aimed at advanced users, who may wish to experiment a little or who find themselves in unusual situations like the necessity to write software on platforms that do not have threads (like classic Unix or some embedded systems). Basically, you will rarely (if ever) need to write your own event processing.

It is possible to compile the *YAMI* library without threading. Then, most of the library's functionality is gone and it may seem at first that it is not possible to use *YAMI* at all. But it is.

In order to correctly process connections coming from remote *Agents*, the local *Agent* performs a loop, where new connections are managed and where everything is made. In fact, every thread in the *Agent* performs its own loop, but the only really necessary is the loop in the receiver module. There are two situations where there may be no separate receiver module:

- The target platform does not support threading.
- The programmer wants to have ultimate control over *when* the event (the incoming connection) is processed.

In both situations the programmer needs to perform the loop by himself, usually in the main part of his program – this applies both to client and server. It is also possible to combine this *explicit event processing* with other loops that may appear in the crucial part of your programs, like in windows-based GUI applications.

The example programs `clientloop.c` and `serverloop.c` are similar to the examples presented so far, but their *Agents* are prepared to work without separate threads and the important event processing is triggered explicitly in the application code. Previously, the event processing was hidden in the receiver module which was running by itself.

The interesting thing to note is that the explicit event processing results in best performance, since everything related to message processing is done in the context of a single thread and most (or all) of the synchronization overhead is avoided. Nice, isn't it?. Of course, this comes at the cost of increased complexity, as usual.

Last few comments:

- Always check error codes. It will considerably save your debugging time. This tutorial does not contain error handling only for brevity.
- Read the *YAMI* Core Library Reference.

Enjoy! And, of course, tell your friends about *YAMI*.