

Exploring Lua for Concurrent Programming

Alexandre Skyrme

(Pontifical Catholic University of Rio de Janeiro (PUC-Rio)
Rio de Janeiro, Brazil
askyrme@inf.puc-rio.br)

Noemi Rodriguez

(Pontifical Catholic University of Rio de Janeiro (PUC-Rio)
National Education and Research Network (RNP)
Rio de Janeiro, Brazil
noemi@inf.puc-rio.br)

Roberto Ierusalimschy

(Pontifical Catholic University of Rio de Janeiro (PUC-Rio)
Rio de Janeiro, Brazil
roberto@inf.puc-rio.br)

Abstract: The popularization of multi-core processors and of technologies such as hyper-threading demonstrates a fundamental change in the way processors have been evolving and also increases interest in concurrent programming, particularly as a means to improve software performance. However, concurrent programming is still considered complex, mostly due to difficulties in using the available programming models, which have been subject to recurring criticism. The increased interest in concurrency and the lack of proper models to support it stimulates the development of proposals aimed at providing alternative models. In this paper, we work with some of Lua's facilities for concurrent programming, such as coroutines, threads, and message passing. We also describe the main characteristics of the explored model and present a library developed to implement it, along with results of a performance evaluation.

Key Words: Lua, preemptive multithreading, non-preemptive multithreading, concurrency, luaproc, message passing

Category: D.1.3, D.3.2, D.3.3

1 Introduction

Regardless of its growing importance, concurrent programming is still mostly based on dated models. Constructions like semaphores [Dijkstra 1983], conditional critical regions [Hoare 1972], guards [Dijkstra 1975], and monitors [Hansen 1974, Hoare 1973] were all originally designed for operating systems and are admittedly complex for higher-level programming. Moreover, they do not scale well for massive concurrency. This scenario has stimulated the proposal of alternative models and constructions for concurrent programming, such as Erlang [Armstrong 1996], Polyphonic C# [Benton, Cardelli and Fournet 2002], Sequential

Object Monitors [Caromel, Mateus and Tanter 2004], and the Concurrency and Coordination Runtime [Chrysanthakopoulos and Singh 2005].

Since 2003 the Lua programming language [Ierusalimsky et al. 1996, Ierusalimsky et al. 2006, Ierusalimsky et al. 2007] features coroutines, which enables collaborative multithreading. However, a common criticism of coroutines is that they cannot explore hardware parallelism, such as provided by multi-core processors. In 2006, Ierusalimsky [Ierusalimsky 2006] proposed the use of multiple independent states in Lua to implement *Lua processes*, based on some form of message passing. In this paper we advance that proposal, building a complete library for concurrent programming in Lua based on message passing over channels. As we will see, the resulting library showed encouraging performance results even when running hundreds of thousands of simultaneous processes.

The rest of this paper is organized as follows. In section 2 we point out some of the downsides of multithreading and present the model which we chose to explore for concurrent programming in Lua. In section 3 we describe how we implemented this model and in section 4 we present some results of a performance evaluation of the implementation. Finally, in section 5, we draw some conclusions.

2 Concurrent Programming in Lua

Programming with preemptive multithreading and shared memory demands synchronization constructions to ensure mutual exclusion and conditional synchronization [Andrews and Schneider 1983]. Unfortunately, the synchronization burden, the difficulty to debug code, and the lack of determinism during execution make development with preemptive multithreading and shared memory admittedly complex [Lee 2006].

Moreover, as argued by Ousterhout [Ousterhout 1996], the criticism of multithreading is not limited to development complexity. It is often difficult to obtain good performance when using preemptive multithreading with shared memory. A too coarse locking reduces the opportunities for concurrency, while a fine-grained locking may add too much overhead to the program. Due to these difficulties, many standard libraries are not *thread-safe*, that is, cannot ensure proper behavior of their functions during simultaneous execution by multiple threads, hampering software from exploring this kind of multithreading. Problems with performance are greatly increased when scaling up to massive multithreading.

These issues have encouraged the development of alternative models for concurrent programming. In this work we explore a model based on execution threads with no shared memory, which uses message passing for synchronization and communication. We implemented this concurrency model for Lua through a library called `luaproc`. Next, we describe the model details and the API provided by this library.

Because they have independent resources, we call each thread in the library a *Lua process*. We create Lua processes through calls to the `luaproc.newproc` function. As user threads, Lua processes are entities scheduled exclusively through a scheduler which runs in user space, without direct relation to operating system processes or other kernel scheduled entities.

Communication between Lua processes occurs exclusively through message passing. On the one hand, communication with message passing can be slower when compared to shared memory. On the other hand, the lack of shared memory avoids the performance and complexity penalties associated to shared-memory synchronization primitives. Besides, programs can use the same communication model for processes within the same machine and for processes in a distributed environment.

As their own names imply, the `luaproc.send` function sends messages and the `luaproc.receive` function receives messages. Message addressing is based on *channels*. Channels must be explicitly created by the `luaproc.newchannel` function and destroyed by the `luaproc.delchannel` function. A channel is an entity on its own, without any direct relation to Lua processes. Each channel is named by a string, which must be specified as a parameter to the `luaproc.newchannel` function. Each process may send to and receive from any channel, as long as the process knows the channel name. Thus, it suffices to know a channel name in order to use it.

Each message carries a tuple of atomic Lua values: strings, number, or booleans. More complex types must be encoded in some form. For instance, it is easy in Lua to serialize data [Ierusalimsky 2006], that is to convert it into a stream of bytes or characters, in order to save it in a file, send it through a network connection or, in this case, send it in a message. Structured values can easily be encoded as a piece of Lua code (as a string) that, when executed, reconstructs that value in the receiver.

The `luaproc.send` operation is blocking: it returns only after another Lua process has received its message on the targeted channel or if the channel does not exist. Otherwise the sending Lua process is blocked until one of these two conditions happen.

The `luaproc.receive` function, on the other hand, can be either blocking or non-blocking, depending on a parameter. A blocking receive behaves similarly to a blocking send: it only returns after matching with a send operation on that channel, or if the channel does not exist. The non-blocking receive operation, in contrast, always returns immediately; its result indicates whether it got any message.

The reason we opted for blocking on send operations is that this provides a simpler, more deterministic, programming model. When a call to `luaproc.receive` returns successfully, it is possible to assert that the message was received. Addi-

tionally, non-blocking sends increase implementation complexity. Of course, the programmer can still send messages without blocking the main execution flow. One easy way to do that is to create a new Lua process with the sole purpose of sending a message. Because the creation of Lua processes is a non-blocking operation, control is immediately returned to the creator process and any occasional communication block affects only the newly created Lua process. Another option is to create a “queue process” to enqueue messages to a channel.

3 Model Implementation

In its standard configuration, Lua includes concurrent programming support through the use of coroutines. Each coroutine represents a different execution flow in user space. Execution control relies on a cooperative model and can be accomplished through calls to the `coroutine.yield` and `coroutine.resume` functions. Calls to the `coroutine.yield` function suspend the coroutine’s execution, while calls to the `coroutine.resume` resume it. Once a coroutine starts running, it runs until it finishes or yields.

The API that Lua offers to C includes a function to create coroutines, as well as functions to suspend and resume their execution. This facility, allied to the flexibility offered by the API for interaction with the Lua interpreter from C code and to the dissociation between coroutines and kernel threads, makes Lua particularly well suited for the exploration of our chosen model for concurrent programming.

Another important facility offered by Lua is the possibility of multiple Lua states. The entire API that Lua offers to C operates over an abstract type called a *Lua state*. By creating multiple Lua states, a C program can have multiple Lua programs that are completely independent.

Our library, `luaproc`, uses Lua states and coroutines to implement Lua processes. Each process runs as an exclusive coroutine inside its own Lua state. These processes are run by *workers*, which are kernel threads implemented with the POSIX Threads library (`pthread`) [IEEE 1995]. There is no fixed relationship between workers and Lua processes. Each worker repeatedly gets a process from the ready queue and runs it until it finishes or blocks. Even though we use kernel threads, there is no memory shared among Lua processes, because each has its own Lua state.

The following sub-sections present a more detailed description of our library’s implementation and characteristics.

3.1 Lua Processes

Using Lua code from within C code is normally preceded by the creation of a Lua state, represented in C by a variable of type `lua_State`. A Lua state defines

the interpreter's state and keeps track of functions and global variables, among other information related to the interpreter.

Once Lua code has been loaded in a Lua state, it is possible to control its execution through functions provided by the C API for Lua. Control takes place as if the Lua code was executed as a coroutine. Therefore, even if the Lua code does not include explicit calls to Lua's standard coroutine handling functions, it is possible to suspend and resume its execution through C functions. This feature is essential to allow control over Lua processes execution.

Each Lua process is comprised by an independent Lua state, where the process code is loaded during process creation. The independence between Lua states ensures the lack of shared memory between Lua processes and helps to enforce message passing as a means for interprocess communication. The remaining structure used to implement Lua processes is compact and has few members other than the process Lua state. Among relevant structure members are the process execution state (idle, ready, blocked or finished) and the number of arguments that must be used when resuming its execution in case it is blocked. No unique process identifier (PID) is included since there is no fixed relation between workers and processes.

Even though the creation of a Lua state is a cheap operation, loading all standard Lua libraries can take more than ten times the time required to create a state [Ierusalimschy 2006]. Thus, to reduce the cost of creating Lua processes, only the basic standard library and our own library are automatically loaded into each new Lua process. The remaining standard libraries (io, os, table, string, math, and debug) are pre-registered and can be loaded with a standard call to Lua's `require` function.

Our library also offers a facility to recycle Lua processes, which is optionally activated through a call to the `luaproc.recycle` function. Recycling consists in reusing states from finished Lua processes to execute new processes. Instead of being destroyed after finishing its execution, a state can be stored for reuse. Creation of a Lua process can then be done by loading new Lua code in a recycled state, thus eliminating the costs of creating a new state and loading libraries.

3.2 Scheduler

The scheduler is automatically initialized when our concurrent programming library is loaded. During its initialization, which occurs in the context of the operating system thread responsible for executing the code that loads our library, a worker is created.

The scheduler manages a single ready queue (FIFO), which holds Lua processes ready for execution. The scheduler itself is responsible for adding newly created Lua processes to the end of the ready queue. Workers execute the Lua code associated with each Lua process.

Workers are simply kernel threads, managed with the POSIX Threads library (pthreads), which perform the following cycle: it retrieves the first Lua process from the ready queue; executes the Lua code associated with the process until it finishes, blocks or yields; and takes appropriate measures depending on execution outcome. Creation and destruction of workers in execution time is supported through the API functions `luaproc.createworker` and `luaproc.destroyworker`.

If the execution of a Lua process ends because the Lua code related to the process has finished normally, the worker closes the corresponding Lua state and destroys the process. If, during the execution of a Lua process, a call is made to the standard Lua function `coroutine.yield`, the worker simply reinserts the process at the end of the ready queue. This suspends the process execution and allows other processes to execute, which is the expected behavior of a yield. If the execution of a Lua process results in an unexpected error, the worker prints an error message, closes the corresponding Lua state and destroys the process.

Since there is only a single ready queue, all workers must get Lua processes from the same queue. This implies that shared memory synchronization primitives had to be used to serialize access and manipulation of the queue. To that matter, conditional variables and mutual exclusion were used, as they are both supported by the POSIX Threads library (pthreads).

3.3 Inter-process Communication

Lua uses a virtual stack to pass values to and from C. Each element in this stack represents a Lua value. Calls from Lua to functions implemented in C use the virtual stack to pass function arguments. Likewise, these C functions use the virtual stack to pass results back to Lua. Therefore, passing messages in our library simply implies copying data from the sender's virtual stack to the receiver's virtual stack.

3.4 Blocking Strategy

In our library, a Lua process can only have its execution blocked in two distinct situations:

1. when it calls the blocking receive function with a channel where there are no processes waiting to send, that is, when an attempt to receive a message occurs without a previous corresponding attempt to send to the same channel;
2. when it calls the send function with a channel where there are no processes waiting to receive, that is, when an attempt to send a message occurs without a previous corresponding attempt to receive from the same channel.

When a Lua process blocks, the worker adds it to the corresponding channel's queue and gets another process from the ready queue in order to run it. A blocked Lua process is unblocked only if there is a matching call on the same channel or if the channel where it is blocked is destroyed. When such a matching call happens, the same worker that is executing the process that made the call removes the blocked process from the channel queue, copies message data between virtual stacks, and places the unblocked process at the end of the ready queue.

To keep track of Lua processes that are blocked trying to communicate, each channel has two distinct queues (FIFO): one holds processes blocked when trying to send messages to the channel and another holds processes blocked when trying to receive messages from the channel. At most one of these queues will not be empty at any given time, otherwise the processes from each queue could match.

3.5 A Sample Application

In this section we present, in listing 1, the source code of a sample “hello world” application developed with our library.

Listing 1: A simple ‘hello world’ application with Lua processes.

```

-- load our concurrent programming library
require "luaproc"

-- create an additional worker
luaproc.createworker()

-- create a new lua process
luaproc.newproc( [=[
  -- create a new channel
  luaproc.newchannel( "achannel" )
  -- create a new lua process
  luaproc.newproc( [=[
    -- send a message to the channel
    luaproc.send( "achannel", "hello world" )
  ]=] )
  -- create a new lua process
  luaproc.newproc( [=[
    -- receive a message from the channel
    msg = luaproc.receive( "achannel" )
    -- print the received message
    print( msg )
  ]=] )
]=] )

-- wait until all lua processes
-- have finished before exiting
luaproc.exit()

```

As we can see, the program begins by loading our library with the standard Lua require function. Then, it creates an additional worker and a main Lua

process that will hold the remaining of our application. This main Lua process creates a channel and two additional Lua processes. While one of these processes sends a message on the channel, the other one receives the message and then prints it. We ensure that our application will not exit before all Lua processes have completed their execution by calling the `luaproc.exit` function, which simply prevents workers from exiting while there are unfinished Lua processes.

4 Performance Evaluation

In this section we describe some experiments we made to evaluate the performance of `luaproc`. All tests, unless stated otherwise, were conducted on a computer running the Linux operating system, with an AMD Athlon 64 X2 dual-core 3600+ processor with 3 GB of RAM. The chosen Linux distribution was Ubuntu 7.10 (Gutsy Gibbon), with standard kernel 2.6.22-14-generic #1 SMP and Native POSIX Threads library (NPTL) 2.6.1. All tests used two workers in order to exploit both processor cores and stimulate parallelism.

The tests ran on an unprivileged user account on the operating system. Each test was executed at least three times and the results presented in this section correspond to the values' arithmetic mean. Execution times were measured with Linux's Bourne Again shell (bash) `time` command.

We also carried out some comparative tests to evaluate our library against Erlang [Armstrong 2007]. Despite syntax heterogeneity and implementation differences, notably Erlang's built-in support for most of the functionalities we implement through a library, these tests represent an important benchmark. Moreover, further analysis of their results could result in future improvements to our library.

Erlang offers three different execution modes: interpreted code (`escript`), compiled code with symmetric multiprocessing (SMP) support enabled (`erl -smp`), and compiled code with SMP support disabled (`erl`). In all our tests there was just a slight variation in execution time between compiled code with SMP support enabled and compiled code with SMP support disabled. In fact, our tests consistently showed slightly worse (higher) execution times when SMP support was enabled. Furthermore, SMP support is disabled by default for compiled code and not supported for interpreted code. For these reasons, we choose only to present times for interpreted code and compiled code with SMP support disabled.

Further comparative tests were also carried out in order to evaluate our library against a traditional kernel multithreading with shared memory model by using the POSIX Threads library (pthreads) [Skyrme 2008]. We believe this comparison is not as interesting as the one with Erlang, since the pthreads library does not include built-in message passing primitives and is not known to

offer scalability that allows for massive concurrency. Therefore, we opted not to present them in this work.

4.1 Process Creation

In this simple test we measure execution time to create increasing number of Lua processes. First, a main Lua process which will host the remainder of our application's code is created. Then, from within the main Lua process, the same number of Lua processes and communication channels are created, as if each Lua process had its own channel. The spawned Lua processes simply wait for a message from the main Lua process, which is only sent after all of them have been spawned, and then finish their execution.

We reproduced a similar test with Erlang. Just as in the Lua processes test, a certain number of Erlang processes are created and wait for a message, sent by a main process, before finishing their execution. The main difference is that in Erlang there is no need (nor support) to create communication channels, since messages are addressed using process identifiers (PIDs).

Figure 1 shows the total execution times for creating increasing quantities of Lua processes, along with the total execution times for creating, both using interpreted code and compiled code, Erlang processes. Erlang's interpreter (`escript`) limits process creation to 30,000 processes, which explains why the corresponding line in the figure reaches an upper bound at that point.

As we can observe in the figure, execution time increased almost linearly with the number of processes, both in Lua and in Erlang.

4.2 Memory Usage

In this test we measure the memory usage caused by process creation. Like in the previous test, we first create a main Lua process and then, from within that process, we create the same number of communication channels and Lua processes which wait for a message that is only sent after all processes have been spawned. However, for this test, we introduced delays immediately before and immediately after creating channels and processes, in order to allow for external memory usage measurement with Linux's `pmap` command, which maps virtual memory usage per process.

We ran an analogous test with Erlang. Also like in the previous test, we created a certain number of Erlang processes which waited for a message that was only sent after all processes had been created and we did not create or use communication channels. For this test, though, we introduced delays immediately before and immediately after creating processes, in order to measure memory usage by the same means described previously.




Figure 1: Execution times to create n processes in Lua and Erlang.

Initial memory usage, before process and communication channel creation in Lua, was almost constant and around 33.15MB. In Erlang, on the other hand, initial memory usage before process creation was also almost constant but this time around 25.54MB for interpreted code and 23.53MB for compiled code. Figure 2 shows the total memory usage after the creation of Lua processes and communication channels, along with the total memory usage after the creation of Erlang processes, both using interpreted code and compiled code. Once again, the line in the figure which corresponds to Erlang’s interpreter (`escript`) reaches an upper bound at 30,000 processes due to an interpreter limitation.

As the figure shows, memory usage increased almost linearly with the number of processes in Lua and linearly with the number of processes in Erlang. Regardless of Erlang’s lower memory consumption, the test results confirms our library seems suitable for massive concurrency and offers considerable improvement, concerning memory usage, over traditional kernel based multithreading with shared memory.

4.3 Communication

Message passing is the intended way for Lua processes to communicate and synchronize, therefore it is important to evaluate how it performs. In this test we sequentially send and receive messages of different sizes and measure execution time. First, the message contents are read from a file composed of copies of the

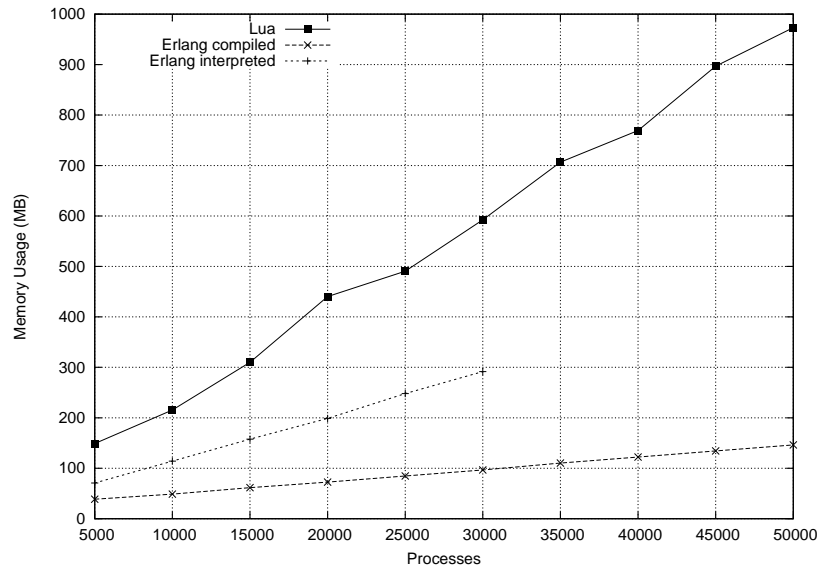


Figure 2: Memory usage measured immediately after creating n processes in Lua and Erlang.

same string separated by newlines. Then, a main Lua process that will host the remainder of our application's code is created. Next, from within the main Lua process, a communication channel is created and a new Lua process, whose sole purpose is to receive messages, is spawned. Finally, the main Lua process sends the same message sequentially, 1,000 times, to the second Lua process.

We conducted a similar test using Erlang. Just as in the previous tests, the main difference between our Lua code and our Erlang code was the lack of need to create communication channels in the later. Apart from that, Erlang code also differs slightly since a few additional messages must be sent in order to inform the receiver process identifier (PID) to the sender process and to ensure proper synchronization.

Figure 3 shows the total execution times for sending and receiving messages of increasing sizes using our library and Erlang. Once again, we present both the results for interpreted and compiled Erlang code.

As we can see in the figure, our library presented good communication performance, with execution times below 0.1s to send messages with up to 10,000 bytes. Erlang, in turn, presented better performance when interpreted, rather than compiled. It presented almost constant execution times when compiled, which suggests it relies on an $O(1)$ operation to perform message passing, such as copying a pointer that points to a shared memory address that holds message data.

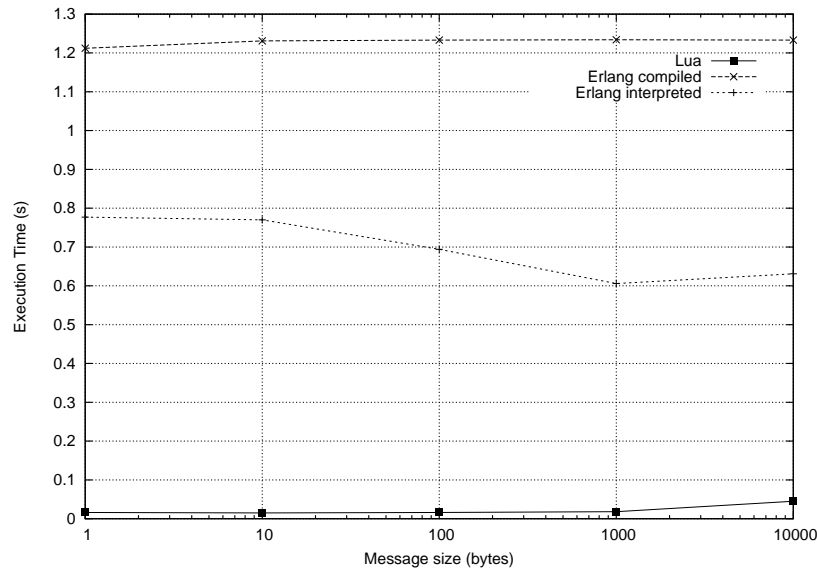


Figure 3: Execution times for sequentially sending 1,000 messages of increasing sizes.

4.4 Process Recycling

In this test, we evaluate how valuable the state recycling feature described in section 3 can be, if used under the right circumstances. This time, we created a fixed number of Lua processes sequentially and had them print a simple message to standard output before exiting. No inter-process communication was used and thus no communication channels were created. The lower individual execution time allowed for better evaluation of the process recycling feature.

We changed the recycled process limit and measured the total execution time for creating and running all the Lua processes. The recycled process limit simply determines the maximum number of Lua states from finished Lua processes which are stored for further recycling when new processes are created. If this limit is set to zero, the recycling feature is disabled and no Lua states from finished Lua processes are kept. If it is set to n , at any given time there will be at most n stored Lua states from finished Lua processes.

An altered version of the library was used for this test just to allow for displaying how many processes were created with recycled Lua states. Figure 4 shows the total execution times for creating 100,000 Lua processes using different recycle limits and redirecting standard output to the null device, along with recycle counts.

As the figure shows, the process recycling feature can offer significant per-

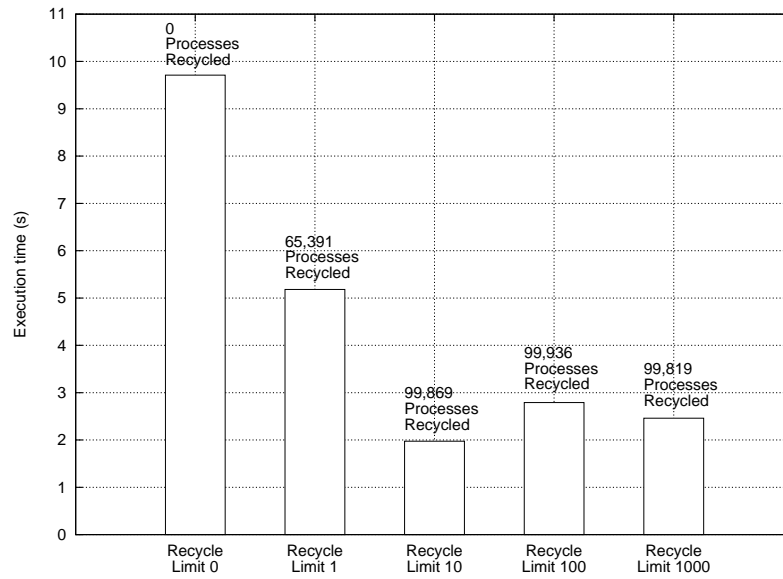


Figure 4: Execution times to create 100,000 processes in Lua using at most n recycled processes.

formance improvements under some circumstances, such as the reduction by almost five times in total execution time observed in this test when using a recycled process limit of just ten processes. However, results also indicate total execution time can cease to decrease, and even increase, after certain recycled process limits. This demonstrates proper care should be taken when using this feature, so the performance cost of recycling processes does not exceed the expected performance improvement.

4.5 Parallelism

The use of multiple workers in multi-processed environments allows for parallel execution of Lua processes. In this test we explore our library's parallelization potential by implementing a parallel string search application. We also implemented a serial version of the same application, using only standard Lua libraries, to evaluate if there is any significant performance cost when our library is used.

The parallel version of the application is divided in three modules. The first module is responsible for initializing the application: it creates workers and communication channels, spawns a coordinator Lua process and several searcher Lua processes, and then sends messages to the coordinator Lua process with the names of the file that holds the patterns and the target to be searched.

The second module is responsible for coordinating job distribution and for centralizing results. It reads the patterns from a file, sends them to the searchers and then starts to progressively distribute target file names to searchers. It also receives results from searchers and notifies them when all target files have been searched.

The third module is the searcher, that is, it is responsible for searching target files for patterns. Each searcher receives a single file name at a time and only sends back its results to the coordinator after processing the whole file. The results sent are composed by the lines of the target file that matched any of the patterns.

This test was exceptionally carried out on a computer with four AMD Opteron dual-core 2.2 GHz processors, for a total of eight processor cores, and 32 GB of RAM. Its operating system was also Linux, but this time with the CentOS 5.1 distribution, standard kernel 2.6.18-53.1.6.el5xen #1 SMP and Native POSIX Threads library (NPTL) 2.5.

Initially, six workers were used to run the parallel version of the application in order to stimulate parallelism and reduce concurrency in the execution of one coordinator and five searcher Lua processes. Next, still using the parallel version of the application, just a single worker was used, in order to allow for a more balanced comparison with the serial version. The pattern file used throughout the test was the same and it contained 25 lines, with one string per line. The target files were copies of a single file, which included 6,605,423 lines and 2,147,483,849 bytes (around 2 GB). Results are shown in figure 5.

The results indicate that exploitation of parallelism on multi-processed environments, as expected, can result in proportional reductions in execution time. As we can observe, when using the serial version of the application or the parallel version with a single worker (kernel thread) execution time increased almost linearly with the number of target files. On the other hand, when the parallel version runs with six workers, execution times for one or five target files was almost the same, which strongly suggests that while one worker acted as the coordinator, the other five workers acted as searchers and processed the five target files in parallel. Still regarding the parallel version with six workers, it is worth noticing that, once again as expected, execution time increased linearly when the number of target files increased from five to ten.

Finally, results also show an almost insignificant difference in execution times between the serial version of the application, which uses only standard Lua libraries, and the parallel version, which uses our library, when it ran with a single worker.

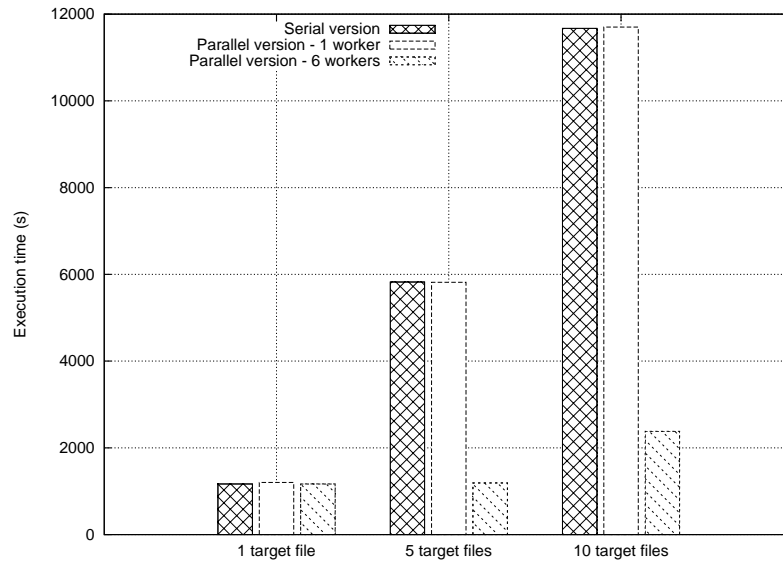


Figure 5: Execution times for serial and parallel string search.

5 Conclusion

In this work we explored an alternative model for concurrent programming in Lua. The model is characterized by using message passing, as opposed to shared memory, as the only inter-process communication method. Its implementation uses kernel threads as workers so that multiple processes can run in parallel.

The lack of shared memory eliminates the need to control access to data shared among execution flows and critical regions, which simplifies development and reduces the probability of inconsistencies which can result in data corruption and execution failures. The predictability of blocked communication, in turn, facilitates debugging and increases determinism of execution flows.

The Lua programming language, despite not being specifically developed for concurrent programming, demonstrated enough flexibility to allow satisfactory implementation of the chosen model for concurrent programming. Additionally, it provided adequate performance and scalability to our library, as can be observed through the results presented in section 4.

Although our library is intended to be used locally only, that is, in individual computers, it is easy to extend it to support execution of Lua processes in a distributed environment. We have already successfully developed and experienced with a very simple client-server application that uses LuaSocket [Nehab 2007], an extension library that adds network support to Lua, to allow for remote creation

of Lua processes.

The use of the POSIX Threads library (pthreads) as a means to benefit from kernel threads allowed for the exploitation of parallelism intermediated by the underlying operating system. Nevertheless, paradoxically, it also resulted in a significant increase in development complexity, mostly due to the need to handle typical obstacles related to using preemptive multithreading with shared memory.

The difficulties we experienced while developing the library to implement the chosen model confirm the criticism of preemptive multithreading with shared memory and reinforce the necessity of newer approaches to concurrent programming. The limitations of preemptive multithreading with shared memory, in particular the complexity in development, create difficulties even when it is used just as a building block to structure alternative solutions.

This work does not exhaust the investigation of the chosen model for concurrent programming in Lua, nor the exploration of new alternatives for concurrent programming. Our library could be further improved by new functionalities and it could be further evaluated by development of more complex, or so-called “real-world”, applications combined with a more extensive performance evaluation. In addition, the usability of our library, which we intuitively believe to be better than other libraries, still lacks proper testing. Nevertheless, the results presented in this work represent an important step towards allowing other contributing efforts to be undertaken.

References

- [Andrews and Schneider 1983] Andrews, G. R., Schneider, F. B.: “Concepts and Notations for Concurrent Programming”; *ACM Comput. Surv.*, 15, 1 (1983), 3–43.
- [Armstrong 1996] Armstrong, J.: “Erlang - a Survey of the Language and its Industrial Applications”; *INAP'96 — The 9th Exhibitions and Symposium on Industrial Applications of Prolog*, Hino, Tokyo, Japan (1996), 16–18.
- [Armstrong 2007] Armstrong, J.: “Programming Erlang”; *Pragmatic Bookshelf* (2007), ISBN 193435600X.
- [Benton, Cardelli and Fournet 2002] Benton, N., Cardelli, L., and Fournet, C.: “Modern Concurrency Abstractions for C#”; *ECOOP '02: Proceedings of the 16th European Conference on Object-Oriented Programming*, Springer-Verlag (2002), ISBN 3-540-43759-2, 415–440.
- [Caromel, Mateu and Tanter 2004] Caromel, D., Mateu, L., and Tanter, E.: “Sequential Object Monitors”; *ECOOP 2004 – Object-Oriented Programming*, 18th European Conference, Springer-Verlag, 3086, 316–340.
- [Chrysanthakopoulos and Singh 2005] Chrysanthakopoulos, G., and Singh, S.: “An Asynchronous Messaging Library for C#”; *Synchronization and Concurrency in Object-Oriented Languages (SCOOL)*, *OOPSLA 2005 Workshop*, San Diego, California, USA.
- [Dijkstra 1975] Dijkstra, E. W.: “Guarded commands, nondeterminacy and formal derivation of programs”; *Commun. ACM* 18, 8 (1975), 453–457.
- [Dijkstra 1983] Dijkstra, E. W.: “The structure of THE - multiprogramming system”; *Commun. ACM* 26, 1 (1983), 49–52.

- [Hansen 1974] Hansen, P. B.: “A Programming Methodology for Operating System Design”; IFIP Congress (1974), 394–397.
- [Hoare 1972] Hoare, C. A. R.: “Towards a theory of parallel programming”; Operating System Techniques (1972), Academic Press, 61–71.
- [Hoare 1973] Hoare, C. A. R.: “Monitors: an operating system structuring concept”; Stanford University, Stanford, CA, USA (1973).
- [IEEE 1995] IEEE: “1003.1c-1995: Information Technology - Portable Operating System Interface (POSIX) - System Application Program Interface (API) Amendment 2: Threads Extension (C Language)”; IEEE Computer Society Press (1995).
- [Ierusalimschy et al. 1996] Ierusalimschy, R., Figueiredo, L. H., and Celes, W.: “Lua - an extensible extension language”; Software: Practice and Experience, 26, 6 (1996), 635–652.
- [Ierusalimschy et al. 2006] Ierusalimschy, R., Figueiredo, L. H., and Celes, W.: “Lua 5.1 Reference Manual”; Lua.Org (2006), ISBN 8590379833.
- [Ierusalimschy et al. 2007] Ierusalimschy, R., Figueiredo, L. H., and Celes, W.: “The Evolution of Lua”; Third ACM SIGPLAN Conference on History of Programming Languages, San Diego (Jun 2007), 2-1–2-26.
- [Ierusalimschy 2006] Ierusalimschy, R.: “Programming in Lua”; Lua.Org (2006), Second Edition, ISBN 8590379825.
- [Lee 2006] Lee, E. A.: “The Problem with Threads”; EECS Department, University of California, Berkeley (2006), UCB/EECS-2006-1
<http://www.eecs.berkeley.edu/Pubs/TechRpts/2006/EECS-2006-1.html>.
- [Nehab 2007] Nehab, D.: “LuaSocket: Network support for the Lua language”;
<http://www.tecgraf.puc-rio.br/luasocket>.
- [Ousterhout 1996] Ousterhout, J.: “Why Threads Are a Bad Idea (for most purposes)”; Presentation given at the 1996 USENIX Annual Technical Conference, January.
- [Skyrme 2008] Skyrme, A.: “An Alternative Model for Concurrent Programming in Lua”; Master’s thesis (2008), Pontifical Catholic University of Rio de Janeiro (PUC-Rio).