Table of Contents

Integrating Kernel Activations in a Multithreaded Runtime System on top of Linux . . . 1
Vincent Danjean, Raymond Namyst, Robert D. Russell
1 Kernel Support for User Level Thread Scheduler . . . 1
1.1 The Marcel-SMP "Mixed" Thread Scheduler . . . 2
1.2 Better Support: Kernel Activations . . . 2
2 Marcel on Top of Linux Activations . . . 3
2.1 Linux Activations . . . 4
2.2 How does it work? . . . 4
2.3 Extensions to the original Activation proposal . . . 5
On the Number of simultaneous Activations . . . 5
Optimizing the number of upcalls . . . 5
2.4 Modifications made within Marcel . . . 6
Global pool of running threads . . . 6
Locks and Preemption in Marcel . . . 6
3 Performance and Evaluation . . . 7
3.1 The test program . . . 7
3.2 Performance . . . 7
Basic version . . . 7
Version with I/O operations . . . 8
Version with substantial computations . . . 8
4 Conclusion . . . 8

Integrating Kernel Activations in a Multithreaded Runtime System on top of Linux

Vincent Danjean[1], Raymond Namyst[1], and Robert D. Russell[2]

[1] Laboratoire de l'Informatique du Parallélisme
École normale supérieure de Lyon
46, Allée d'Italie
F-69364 Lyon Cedex 07, France
{Vincent.Danjean, Raymond.Namyst}@ens-lyon.fr
[2] Computer Science Department
Kingsbury Hall
University of New Hampshire
Durham, NH 03824, USA
rdr@unh.edu

Abstract. Clusters of SMP machines are frequently used to perform heavy parallel computations. The concepts of multithreading have proved suitable for exploiting SMP architectures. Generally, the programmer uses a threads library to write this kind of program. Such a library schedules the threads or asks the OS to do it, but both of these approaches have problems. Anderson et al. have introduced another approach which relies on cooperation between the OS scheduler and the user application using *activations* and *upcalls*. We have modified the Linux kernel and adapted the Marcel threads library (from the programming environment PM^2) to use activations. Improved performance was observed and problems caused by blocking system calls were removed.

1 Kernel Support for User Level Thread Scheduler

The increasing popularity of clusters of SMP machines increases the need for multithreaded programming environments able to fully exploit such architectures. Indeed, the threads model naturally helps to make efficient use of all available processors and to overlap I/O operations with computations. Furthermore, threads are often considered as "virtual processors" and are targeted as such by compilers or runtime support systems for portability purposes. However, these runtime systems are built on top of threads libraries that do not all have the same properties, and thus do not provide the same functionalities. Moreover, these properties directly depend on how much control the threads scheduler has over the architecture's resources. There are two principle kinds of threads: user-level and kernel-level, each with its own advantages and inconveniences.

Efficiency is the main advantage of user-level threads libraries, whose scheduler is completely implemented in user space. Most operations on threads (creations, context switches, etc.) can be done without any call to the operating

system. As a result, certain computations utilizing these threads may perform one or two orders of magnitude better than kernel-level threads. Furthermore, user threads are very much more efficient in terms of kernel resource consumption, which means there can often be many more of them per application. Finally, since user-level threads are implemented in user space, they can be tailored to each user's application. The disadvantage is that user-level threads are "ignored" by the OS and thus cannot be scheduled correctly in many cases. For instance, since user threads within the same process cannot be scheduled concurrently on multiple processors, no real parallelism can be achieved within one application. Similarly, when a thread makes a blocking system call (for example, a `read()` on an empty socket), all the threads of that process are blocked.

Obviously, kernel-level threads do not suffer from these drawbacks, since their scheduling is realized within the OS kernel, which handles them the same way it handles processes, except that multiple threads may share the same address space. It is therefore possible on an SMP machine for the kernel to simultaneously assign processors to multiple threads in the same application, thus achieving true parallelism. Furthermore, when one thread makes a blocking system call the kernel can give control to another thread in the same application. However, even if operations such as thread context switching are more efficient than those related to processes, they still require system calls to be performed which makes them at least an order of magnitude slower than comparable operations for user-level threads.

1.1 The Marcel-SMP "Mixed" Thread Scheduler

To try to obtain the best properties of the two kinds of threads, some libraries mix them together: there are a fixed number of kernel threads each running a number of user threads. This approach retains the efficient scheduling of user threads, but is able to take advantage of parallelism between threads on SMP machines. One such library is MARCEL[4], which was developed for use by PM^2[3] (*Parallel Multithreaded Machine*), a distributed multithreaded programming environment. PM^2 is optimized to support an unpredicatably large number of concurrent tasks, each of which may have different resource requirements and execution times. MARCEL delivers good performance by eliminating some features from the POSIX pthreads specification that are not useful for scientific applications (*e.g.*, per-thread signal handling). In addition, it supports multiple optimizations as well as dynamic thread migration across nodes of a homogeneous cluster. MARCEL has been ported to a number of different platforms. It utilizes a fixed number of kernel threads, each managing a pool of user-level threads.

1.2 Better Support: Kernel Activations

Although the two-level version of MARCEL achieves better performance than an earlier user-level-only version, it still suffers from some of the problems discussed earlier. The first problem is that when a user thread makes a blocking system call,

the underlying kernel thread is stopped too. It is possible with a few blocking user threads to block all the kernel threads, thereby blocking the whole application, even if some other user threads are ready to run. Another problem is that even if Marcel can control the scheduling of user-level threads in each pool, it cannot do anything between the different pools. So, if thread A in pool 1 holds a lock and is preempted by the system, then when thread B in another pool wants the lock, it has to wait for the OS to give control back to pool 1 so that thread A can release the lock.

These problems could be avoided if the OS scheduler reported what it was doing to the application. One mechanism to achieve this cooperation is based on the concept of "activations" which was first proposed in an article by Anderson et al.[1] Its authors implemented this mechanism with the FastThread library on the Topaz system. However, this system is no longer running, and the sources were never released. All the terms (*activation*, *upcall*, etc.) used in this paper come from this article.

This mechanism enables the kernel to notify a user-level process whenever it makes a scheduling decision affecting one of the process's threads. This mechanism is implemented as a set of "upcalls" and "downcalls". A traditional system call is a "downcall", from the user-level down into a kernel-level function. The new idea is a corresponding "upcall", from the kernel up into a user-level function. An upcall can pass parameters, just as system calls do. An "activation" is an execution context (*i.e.*, a task control block in the kernel, similar to a kernel-level thread belonging to the process) that the kernel utilizes to make the upcall. The key point is that each time the kernel takes a scheduling action affecting any of an application's threads, the application receives a report of this fact and can take action to (re)schedule the user-level threads under its control.

We have modified the Linux kernel by adding activations and changing the existing kernel scheduler to use upcalls to report some scheduling events to the Marcel scheduler running in user space. Upcalls are mainly used to report that a new activation has been created, that an activation has blocked in a system call, that a previously blocked activation has just been unblocked, or that an activation has been preempted. We have also modified Marcel to utilize this mechanism efficiently, as discussed in the next section.

2 Marcel on Top of Linux Activations

The user-level Marcel threads manager utilizes the new mechanism as follows: Marcel begins by making an `act_new` system call to notify the kernel that it wants to utilize activations. The manager provides parameters that include a vector of entry points for a fixed set of user-level management functions to which the kernel will make upcalls. Whenever the kernel makes a scheduling decision affecting any of this process's activations, such as creating, blocking or unblocking it, the kernel informs the process by choosing one of its activations and using it to make the appropriate upcall, such as `upcall_new`, `upcall_block`, or `upcall_unblock`. In order to guarantee exclusive access to management infor-

mation while executing one of these functions, the kernel maintains an internal mutual exclusion lock that allows only one upcall at a time to be outstanding per process. Therefore, the management function must make an `act_resume` system call to release that lock after making its management decision but before executing application specific code. If the kernel scheduler decides an activation holding this lock should be preempted, the kernel will preempt another activation instead (via `upcall_preempt`) and will simply reschedule the original activation without an upcall. If one activation terminates (by calling `exit()` for example) then the whole process terminates.

2.1 Linux Activations

Our implementation of the activations within the LINUX kernel is closed to the one proposed by Anderson et al. We have a few system calls and the kernel can make five kinds of upcalls.

System call	Description
act_init	give to the kernel the information needed (addresses of the functions for the upcall, of the stack to use, etc.)
act_resume	allow the kernel to make another upcall
act_cntl	used to query/set some parameters about activations
act_send	used to force an upcall in an activation

Upcall	Description
act_new	a new activation is starting
act_block	an activation blocked
act_unblock	an activation unblocked. We have its state, so we can restart its thread when we want.
act_prempt	an activation was preempted. We have its state, so we can restart its thread when we want, instead of the current one for example.
act_restart	we have to make an upcall (probably because a call to act_send), but we have no event to report, so we use this upcall.

2.2 How does it work?

Figure 1 illustrates how MARCEL uses activations to keep both processors on a dual-processor SMP platform actively executing application threads, even when some threads are blocked in the kernel.

At time T1, the kernel creates two activations "A" and "B" and makes an `upcall_new` to each. In each activation, the MARCEL manager will choose a ready application thread and give it control.

At time T2, the application thread running in activation "A" makes a blocking I/O system call. The kernel determines the process to which this activation belongs and creates a third activation, "C", into which it makes an `upcall_new`. In this activation, the MARCEL manager will choose a third application thread

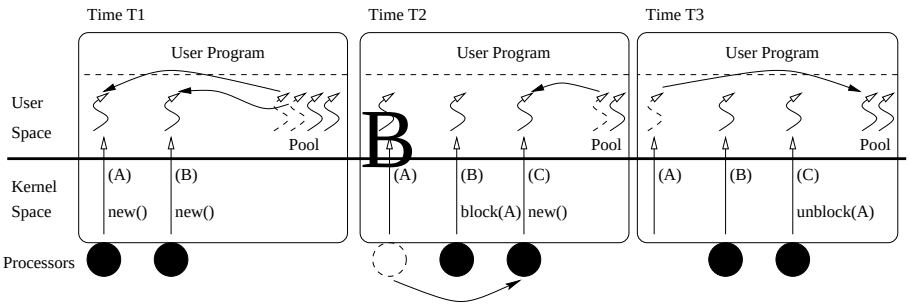

Fig. 1. A Blocking System Call with Activations

and then call `act_resume` to release the mutual exclusion lock. The kernel next chooses one of the activations, say "B", and makes an `upcall_block` to it, providing "A" as the parameter to indicate which activation was just blocked. The MARCEL manager uses this information to keep track of the status of the corresponding application thread.

At time T3, the I/O request completes. The kernel then chooses one of the activations, say "C", and makes an `upcall_unblock` to it, providing "A" as the parameter to indicate which activation was just unblocked. The MARCEL manager now chooses whether to return the application thread previously assigned to "A" to the pool and continue running the application thread already assigned to "C", or vice versa. In either case, activation "A" remains idle until needed by the kernel to make another upcall.

2.3 Extensions to the original Activation proposal

Although this work is mainly based on the *Scheduler Activation* model proposed by Anderson et al, we have proposed a number of improvements which extend the set of supported system calls and increase efficiency in some situations.

On the Number of simultaneous Activations One important point with activations is that the number of concurrently running activations in the application is constant. But this does not mean that the number of activations in the kernel is constant. In the Anderson's implementation, it was. This has the advantage of using constant kernel ressources, but does not allow to handle all the blocking system calls. Indeed, we need to be able to release all the kernel structures before the system call ends when it blocks. Our implementation does not keep constant the number of simultaneous activations within the kernel. This allows us to treat absolutely any system call without limitation. The couterpart is the number of kernel resources that we potentialy use.

Optimizing the number of upcalls Several optimizations have been made to increase the performance of our implementation.

When an activation blocks, we typically need to make two upcalls: one (act_block) to prevent the application that an activation blocked, another (act_new) to launch a new activation. This is now handled by only one upcall to act_new which has another parameter used to tell the application if another activation has blocked.

Another optimisation has been made with the preemption. In the original model, when an activation is preempted, an act_preempt upcall occurs, and an act_new upcall is made when the kernel is ready to restart an activation. Now, the application can tell the kernel at the end of the act_preempt upcall (with a parameter in the system call act_resume) that instead of calling the upcall act_new, it can continue this activation directly.

2.4 Modifications made within Marcel

Surprisingly, the integration of Linux Activations within the Marcel library did require almost no rewriting of the existing code. We rather made a few localized extensions.

Global pool of running threads The major issue that we had to address was related to the "ready-threads" queue. The problem was to opt either for a global pool (as in a user-level version of Marcel) or for a collection of activation-specific local pools (as in the mixed version). We have opted for the global pool implementation because maintaining separate pools introduces a number of problems in this case. In particular, when an activation gets blocked within the kernel, the other activations must retrieve the running threads that were kept in its ready-threads pool. Such a step requires a costly synchronization scheme and the associated overhead may become important in the presence of frequent I/O operations. The drawback of our stragegy is that the global pool may become a bottleneck on a large number of processors.

Locks and Preemption in Marcel Marcel uses an special lock to prevent concurrent access to its internal data structures. Our implementation of activations makes sure that if the kernel preempts the marcel thread which is holding this lock, then it is relaunched immediately (instead of the one running on the activation that receive the `act_preempt` upcall). This allows us to avoid contention situations in the presence of busy waiting threads.

Note that a related problem can occur with the upcall `act_new`. Indeed, when a new activation is created, it may not succeed in acquiring the aforementionned lock. Since it is mandatory to run a regular Marcel thread when calling `act_resume`, the activation must schedule a "dummy" thread. To this end, we have added a pool of preallocated "dummy" threads (together with their stacks) into Marcel.

3 Performance and Evaluation

The new version of Marcel on top of Linux Activations is completely operational, although we did not yet implement all the optimizations we discussed in the previous sections. To investiguate the gain or the overhead generated by activations and upcalls, we have compared the new version of Marcel to the two existing versions (one purely user-level, one mixed two-level) as well as to native Linux kernel-level threads ([2]). The tests were run on Intel Pentium II 450 MHz platform running Linux v2.2.13. On this platform, we ran a microbenchmark program to measure the time taken by an *upcall* from the kernel up to user-space. This test reported an average time of $5\mu s$ per upcall.

3.1 The test program

The tests used to compare these libraries are all based on a common synthetic program. This program basically implements a *divide and conquer* algorithm to compute the sum of the Nth first integers. At each iteration step, two threads are spawned to compute the two resulting sub-intervals concurrently, except if the interval to compute contains one element. The "father" of two threads waits their completion, gets their results, computes the sum and in turn returns it to its father. To sum up, this program generates a tree of threads and involves almost no real computation but a lot of basic thread operations such as creation, destruction and synchronization.

In order to evaluate the different thread libraries in the presence of blocking calls, we have extended the previous program so as to make extensive use of Unix I/O operations. In this case, we have simply replaced all the thread creation calls by a write into a Unix *pipe*. At the other end of the pipe, a dedicated server thread serves simply transforms the corresponding requests into thread creations.

Finally, we also extended the basic version of the program by adding some artificial computation charge intro each thread so that some speedup can be obtained on a multiprocessor platform.

3.2 Performance

Table 1 reports the performance obtained with the three aforementionned program versions for each thread library. The first two programs were run on a monoprocessor machine whereas the last one was run on a dual processor machine.

Basic version The raw version of the divide and conquer program make heavy use of thread creations and synchronizations. On a monoprocessor machine, the user-level Marcel library is obviously the most efficient, as one may expect. On the contrary, the LinuxThread library exhibits poor performance because kernel thread operations are much more inefficient than those related

Table 1. Performance of various thread libraries on a uni-processor

Library	Basic	With I/O	With computations
	Single processor		Dual processor
MARCEL user-level	0.308ms	119.959ms	6932ms
MARCEL mixed two-level	0.435ms	23.241ms	3807ms
MARCEL with activations	0.417ms	10.118ms	3551ms
LINUXTHREAD (kernel-level)	13.319ms	14.916ms	3566ms

to user threads. It is interesting to note that the activation achieves good performance. The difference with the user-level version is due to the Marcel lock acquire/release primitives that are a little more complex in the presence of activations.

Version with I/O operations When many I/O operations are done during the execution of such a program, things change significantly. The most noticeable result is the *huge* duration of the program with the user-level version. It is however not surprising: each time a user thread makes a blocking call, it blocks the entire Unix process until a timer signal forces a preemption and schedules another thread (in this case, every $20ms$). The activation version has the best execution time. The mixed Marcel library does not behave so well because two underlyin kernel threads are needed to handle the blocking calls properly. Thus, it introduces an overhead due to additionnal synchronization and preemption costs.

Version with substantial computations When execution the program on a dual-processor machine, we observe that the activation version has approximately the same execution time than the Marcel "mixed" and LinuxThread versions. It reveals that the activation version is perfectly able to exploit the underlying architecture by using two activations simultaneously within the application. Obviously, the user-level version performs poorly because only one processor is used in this case.

4 Conclusion

This work designed, implemented and tested the use of activations, a new technique to handle thread support in an OS. We wrote a new version of the MARCEL threads library that utilizes activations while preserving the existing user interface, so that existing MARCEL programs still work with this new model. On a uniprocessor the performance is as good as with the best previous version of MARCEL, and on a multiprocessor it is as good as with kernel-level threads.

With the new version of MARCEL, applications can now use OS features that previously caused difficulties. For example, blocking system calls, such as `read()`

or native semaphore operations, are handled correctly and no longer block other threads. Furthermore, since our new library is implemented in user space, we do not need to change the kernel to add new threads features, such as thread migration.

A two-level threads library based on activations seems to be a very attractive way to manage application threads. This work shows that this model is a valid one, in particular for application threads that utilize blocking system calls, which often happens within a communication library for example.

References

1. T. Anderson, B. Bershad, E. Lazowska, and H. Levy. Scheduler Activations: Effective Kernel Support for the User-Level Management of Parallelism. *ACM Transactions on Computer Systems*, 10(1):53–79, February 1992.
2. Xavier Leroy. The LinuxThreads library. http://pauillac.inria.fr/~xleroy/linuxthreads.
3. R. Namyst and J.F. Mehaut. PM2: Parallel Multithreaded Machine. a computing environment for distributed architectures. In *ParCo'95 (PARallel COmputing)*, pages 279–285. Elsevier Science Publishers, Sep 1995.
4. R. Namyst and J.-F. Mhaut. MARCEL *: Une bibliothque de processus lgers.* Laboratoire d'Informatique Fondamentale de Lille, Lille, 1995.