



Center for Embedded Computer Systems
University of California, Irvine

A User-level Thread Library Built on Linux Context Functions for Efficient ESL Simulation

Guantao Liu and Rainer Dömer

Technical Report CECS-13-07
June 6, 2013

Center for Embedded Computer Systems
University of California, Irvine
Irvine, CA 92697-3425, USA
(949) 824-8059

{guantaol, doemer}@uci.edu
<http://www.cecs.uci.edu/>

A User-level Thread Library Built on Linux Context Functions for Efficient ESL Simulation

Guantao Liu and Rainer Dömer

Technical Report CECS-13-07
June 6, 2013

Center for Embedded Computer Systems
University of California, Irvine
Irvine, CA 92697-3425, USA
(949) 824-8059

{guantaol, doemer}@uci.edu
<http://www.cecs.uci.edu>

Abstract

Currently QuickThreads library is widely used in the multi-threaded programs such as Electronic System Level (ESL) simulation. As a user-level thread library, QuickThreads is very efficient in thread manipulation as it operates solely at user level and introduces no operating system overhead. While QuickThreads library utilizes a portable interface to wrap machine-dependent code that performs thread initialization and context switching, it only works on a certain number of specific platforms and architectures. In this report, we propose a new user-level thread library that offers the same features as QuickThreads, but makes use of Linux library functions and therefore is portable to all 32-bit Linux platforms.

Contents

1	Introduction	1
2	Basic Ideas of the ContextThreads Library	2
3	Performance Evaluation of the ContextThreads Library	2
3.1	Platform Architectures and Benchmark Examples	2
3.2	Producer-Consumer Model	5
3.3	Threads with Pure Floating-point Multiplication (TFMUL)	6
4	Conclusion	6
	References	7
A	Benchmark Examples	8
A.1	Producer-Consumer Model	8
A.2	TFMUL Model	13
B	Measured Simulation Times for All Benchmarks and Applications	15
B.1	Simulation Time for Producer-Consumer Model	15
B.2	Simulation Time for TFMUL Model	17

List of Figures

1	Intel Core 2 Quad architecture, Q9650 (μ) [3]	3
2	Intel Xeon architecture, X5650 (ξ) [3]	3
3	Simulation Results for Producer-Consumer Model on μ	4
4	Simulation Results for Producer-Consumer Model on ξ	4
5	Simulation Results for TFMUL Model on μ	5
6	Simulation Results for TFMUL Model on ξ	5

List of Tables

1	Simulation Results for Producer-Consumer Model	4
2	Simulation Results for TFMUL Model	5
3	Producer-Consumer Model on μ	15
4	Producer-Consumer Model on ξ	16
5	TFMUL Model on μ	17
6	TFMUL Model on ξ	17

List of Listings

1	Producer-Consumer Model	8
2	TFMUL Model	13

A User-level Thread Library Built on Linux Context Functions for Efficient ESL Simulation

Guantao Liu and Rainer Dömer

Center for Embedded Computer Systems
University of California, Irvine
Irvine, CA 92697-3425, USA

{guantaol, doemer}@uci.edu
<http://www.cecs.uci.edu>

Abstract

Currently QuickThreads library is widely used in the multi-threaded programs such as Electronic System Level (ESL) simulation. As a user-level thread library, QuickThreads is very efficient in thread manipulation as it operates solely at user level and introduces no operating system overhead. While QuickThreads library utilizes a portable interface to wrap machine-dependent code that performs thread initialization and context switching, it only works on a certain number of specific platforms and architectures. In this report, we propose a new user-level thread library that offers the same features as QuickThreads, but makes use of Linux library functions and therefore is portable to all 32-bit Linux platforms.

1 Introduction

Before the popular use of multiprocessor machines, user-level thread library is widely adopted in multithreading programs to handle multiple concurrent tasks in a time-slice manner. On the single processor, multiple tasks are executed in a time-division multiplexing mode and the context switching between these tasks generally happens so frequently that the users perceive the threads or tasks as running at the same time. Thus, the essential part of the user-level thread library is how to efficiently manage context switching between different user threads in the same process. A well-known example of the user-level threads is QuickThreads [2], which owns great performance in sequential multithreading programs. However, QuickThreads only works right on a limited number of platforms, which restricts the usage of QuickThreads. In the remainder of this technical report, a new

user-level thread library built on Linux context library functions is proposed for the SpecC sequential simulator. Preliminary simulation results indicate that the new thread library has very similar performance to QuickThreads and it is portable to all 32-bit Linux platforms.

2 Basic Ideas of the ContextThreads Library

As a popular user-level thread library, QuickThreads provides a somewhat portable interface to machine-dependent code that performs thread initialization and context switching [2]. To offer more simplicity and flexibility to the thread package, QuickThreads separate the notion (starting and stopping threads) from the thread allocation and scheduling of different queues. In fact, QuickThreads does not manipulate any allocation and run queues. Instead, it only provide a simple mechanism that performs a context switch, and then invokes a client function on behalf of the halted thread. During such a context switching, QuickThreads library will first save the register values of the old thread on to its stack, adjusting the stack pointer as needed, and then jump to the functions of the new thread by loading its stack.

Although QuickThreads library is superior in the thread initialization and context switching, it is only portable to a certain number of platforms and architectures (80386, DEC, VAX family and so on). On other platforms, modern Linux operating systems offer some library functions to achieve the same functionalities. Some examples of these functions are *getcontext*, *setcontext*, *swapcontext* and *makecontext* declared in *ucontext.h* [1]. Specifically, *getcontext* would save the current execution context (register values, program counter and etc.) to a data structure typed *ucontext_t*, *setcontext* would load a *ucontext_t* struct and switch to the specified execution context, *swapcontext* will save the current context and switch to another one and *makecontext* will create a new execution context by defining the thread functions and arguments. To create a new thread, we can first call *getcontext* to retrieve the current context and modify the context by specifying the function and arguments in *makecontext*. To context switch to a new thread, we could just use the *swapcontext* function to stop the current thread and continue executing another. Integrating these functions, a new user-level ContextThreads library is created. Similar to what QuickThreads do, ContextThreads separate the thread execution and scheduling. Changing scheduling policies in the ContextThreads library would be as easy as changing a function pointer in *makecontext*. By utilizing the same ideas in QuickThreads library and the Linux context functions, the new user-level thread library offers high performance in computation as well as a wide portability to all Linux platforms.

3 Performance Evaluation of the ContextThreads Library

3.1 Platform Architectures and Benchmark Examples

To evaluate the performance of ContextThreads, we utilize two SpecC benchmarks to test two different aspects of a thread library: a Producer-Consumer example (Prod-Cons) to evaluate the context switching performance and a parallel benchmark which has intensive thread creation/deletion operations to test the feature of thread initialization (named TFMUL, Threads with pure Floating-point MULTIplication). Both of the benchmarks are running on two 32-bit Linux machines, which have

Intel(R) Core(TM) 2 Quad architecture Q9650 3.0 GHz CPU (named **mu**) and Intel(R) Xeon(R) architecture X5650 2.66 GHz CPU (named **xi**) respectively. Figure 1 and 2 illustrate the architectures of the two processors. The dashed line in the middle of the processor means that the CPU has the hyperthreading feature enabled [3].

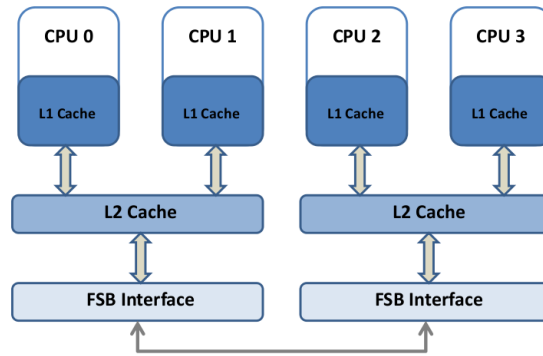


Figure 1: Intel Core 2 Quad architecture, Q9650 (mu) [3]

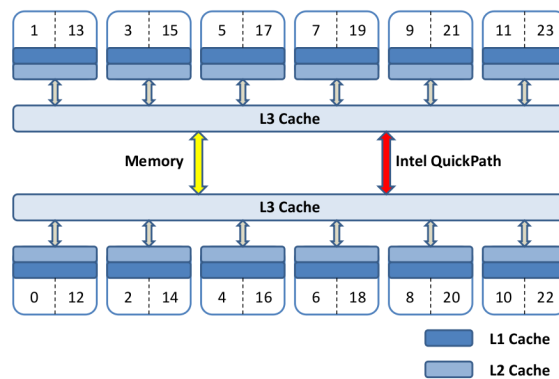


Figure 2: Intel Xeon architecture, X5650 (xi) [3]

To simulate these two examples, we adopt three SpecC sequential simulators which are based on QuickThreads, ContextThreads and PosixThreads. The simulation times for the two benchmarks on **mu** and **xi** are shown in Figure 3 to 6¹. Table 1 and 2 list the data used in these figures.

¹The simulation results in all these figures are picked up from the tables in Appendix B, and they always choose the example which has medium elapsed time

Table 1: Simulation Results for Producer-Consumer Model

Hostname	Usr Time	Sys Time	Elapsed Time	CPU Load	Thread Library
mu	26.85s	0	26.87s	99.00%	QuickThreads
	34.11s	14.22s	48.35s	99.00%	ContextThreads
	84.8s	189.48s	274.38s	99.00%	PosixThreads
xi	22.08s	0	22.14s	99.00%	QuickThreads
	28.75s	9.79s	38.65s	99.00%	ContextThreads
	63.6s	231.25s	295.66s	99.00%	PosixThreads




Figure 3: Simulation Results for Producer-Consumer Model on mu




Figure 4: Simulation Results for Producer-Consumer Model on xi

Table 2: Simulation Results for TFMUL Model

Hostname	Usr Time	Sys Time	Elapsed Time	CPU Load	Thread Library
mu	8.61s	17.67s	26.29s	99.00%	QuickThreads
	11.1s	24.36s	35.48s	99.00%	ContextThreads
	38.22s	169.36s	230.67s	89.00%	PosixThreads
xi	7.38s	11.21s	18.69s	99.00%	QuickThreads
	10.1s	16.72s	26.93s	99.00%	ContextThreads
	37.84s	163.84s	222.02s	90.00%	PosixThreads




Figure 5: Simulation Results for TFMUL Model on mu




Figure 6: Simulation Results for TFMUL Model on xi

3.2 Producer-Consumer Model

The first parallel benchmark, Producer-Consumer model, is a simple example which features intensive context switching. During the whole simulation, the program will create three threads: a

Producer, a Consumer and a Monitor. The Producer instance will repeatedly send data to the Consumer through a double-handshake channel. This communication is wrapped up in a huge loop and the monitor will terminate the whole program when all the communication is done. Hence, this example has a limited amount of parallelism but a heavyweight of thread synchronization. The exact code of the Producer-Consumer model is listed in Appendix A.1.

From Figure 3 and 4, it is easily seen that QuickThreads library has the best performance on both two platforms while the sequential PosixThreads simulator owns the worst performance. ContextThreads library is slightly inferior to QuickThreads and is much better than PosixThreads. The almost zero system time in the QuickThreads simulator indicates that QuickThreads have very low kernel-level overhead and is quite efficient in context switching. ContextThreads library has a small amount of system-level time as the signal mask is saved and restored using the system call *sigprocmask*, introducing some kernel-level overhead to the ContextThreads library. Even so, ContextThreads is still more than 5 times faster than PosixThreads for the Producer-Consumer example. When the benchmark has intensive context switching, PosixThreads spend lots of time in the kernel-level scheduling and synchronization of different threads. From Figure 3 and 4, we can see that the system-level overhead of PosixThreads is more than 10 times larger than that of QuickThreads and ContextThreads. Generally speaking, for the Producer-Consumer Model, the sequential QuickThreads simulator has a speedup of 10 over the sequential PosixThreads simulator and the ContextThreads simulator has a speedup of more than 6 over the PosixThreads simulator.

3.3 Threads with Pure Floating-point Multiplication (TFMUL)

TFMUL model is a highly parallel example that stresses thread creation/deletion. In each thread of TFMUL, it is doing floating-point multiplication and there is no inter-thread communication. Thus, all the child threads in the benchmark can be executing at the same time without any data dependencies. A total of 10,000,000 threads are created in the whole process and it brings a heavy load in thread initialization. A.2 shows the source code of the TFMUL example.

For the TFMUL benchmark, we can draw the same conclusion as the Producer-Consumer example. The heavy load of kernel threads creation and kernel structs manipulation will burden the performance of PosixThreads library. On both **mu** and **xi**, the PosixThreads library is more than 6 times slower than the other two user-level thread libraries. For QuickThreads and ContextThreads, they both have a high efficiency in thread initialization (as indicated by the small system time on Figure 5 and 6) and QuickThreads library still owns the best performance.

4 Conclusion

According to the simulation results for the two benchmarks, we can conclude that ContextThreads library is slightly inferior to QuickThreads as a user-level thread library, but is portable to all 32-bit Linux platforms as it does not depend on any machine-dependent code (ContextThreads fail on LP64-architectures as the function *makecontext* requires additional parameters to be type `int`, but the function call passes pointers. On LP64 architectures, the size of pointer is larger than that of an integer). Therefore, ContextThreads library would be a good substitute of the QuickThreads library

on platforms where QuickThreads are not configurable.

Acknowledgment

The authors thank Professor Demsky, EECS Department, UC Irvine for the initial idea of replacing QuickThreads with ContextThreads.

References

- [1] GNU. *Linux Programmer's Manual*, March 2009.
- [2] David Keppel. Tools and Techniques for Building Fast Portable Threads Packages. Technical Report UWCSE 93-05-06, University of Washington Department of Computer Science and Engineering, May 1993.
- [3] Guantao Liu and Rainer Dömer. Performance Evaluation and Optimization of A Custom Native Linux Threads Library. Technical Report CECS-12-11, Center for Embedded Computer Systems, University of California, Irvine, October 2012.

A Benchmark Examples

A.1 Producer-Consumer Model

Listing 1: Producer-Consumer Model

```
1 // prodcons.sc: simple producer-consumer example
2 // author:      Rainer Doemer, Guantao Liu
3 // 02/14/13 GL modified to test HybridThreads library
4 // 01/09/12 RD modified to test ooo-simulator run-ahead
5
6 #include <stdio.h>
7 #include <stdlib.h>
8 #include <assert.h>
9 #include <sim.sh>
10 #include <sched.h>
11
12 #include <c_typed_double_handshake.sh>
13
14 #define DATA 42
15
16 #define ITERATIONS 5000000
17
18 #define EXIT_ON_HANDSHAKE // to exit on hand-shake
19
20 #define printf nop // eliminate printing messages
21
22 #ifndef USE_FLOATING_POINT
23 #define FDATA 42.5e0
24 #endif
25
26 #ifndef MAXTHREAD
27 #define MAXTHREAD 1
28 #endif
29
30 DEFINE_L_TYPED_SENDER(char, char) // interface i_char_sender
31 DEFINE_L_TYPED_RECEIVER(char, char) // interface i_char_receiver
32 DEFINE_L_TYPED_TRANCEIVER(char, char) // interface i_char_tranceiver
33 DEFINE_C_TYPED_DOUBLE_HANDSHAKE(char, char) // channel c_char_double_handshake
34
35 DEFINE_L_TYPED_SENDER(short, short) // interface i_short_sender
36 DEFINE_L_TYPED_RECEIVER(short, short) // interface i_short_receiver
37 DEFINE_L_TYPED_TRANCEIVER(short, short) // interface i_short_tranceiver
38 DEFINE_C_TYPED_DOUBLE_HANDSHAKE(short, short) // channel c_short_double_handshake
39
40 DEFINE_L_TYPED_SENDER(int, int) // interface i_int_sender
41 DEFINE_L_TYPED_RECEIVER(int, int) // interface i_int_receiver
42 DEFINE_L_TYPED_TRANCEIVER(int, int) // interface i_int_tranceiver
43 DEFINE_C_TYPED_DOUBLE_HANDSHAKE(int, int) // channel c_int_double_handshake
44
45 DEFINE_L_TYPED_SENDER(long, long) // interface i_llong_sender
46 DEFINE_L_TYPED_RECEIVER(long, long) // interface i_llong_receiver
47 DEFINE_L_TYPED_TRANCEIVER(long, long) // interface i_llong_tranceiver
48 DEFINE_C_TYPED_DOUBLE_HANDSHAKE(long, long) // channel c_llong_double_handshake
49
50 #ifndef USE_FLOATING_POINT
51 DEFINE_L_TYPED_SENDER(float, float) // interface i_float_sender
52 DEFINE_L_TYPED_RECEIVER(float, float) // interface i_float_receiver
53 DEFINE_L_TYPED_TRANCEIVER(float, float) // interface i_float_tranceiver
```

```

54 DEFINE_C_TYPED_DOUBLE_HANDSHAKE(float , float)// channel c_float_double_handshake
55
56 DEFINE_L_TYPED_SENDER(double , double) // interface i_double_sender
57 DEFINE_L_TYPED_RECEIVER(double , double) // interface i_double_receiver
58 DEFINE_L_TYPED_TRANCEIVER(double , double) // interface i_double_tranceiver
59 DEFINE_C_TYPED_DOUBLE_HANDSHAKE(double , double)// channel c_double_double_handshake
60
61 DEFINE_L_TYPED_SENDER(ldouble , long double) // interface i_ldouble_sender
62 DEFINE_L_TYPED_RECEIVER(ldouble , long double) // interface i_ldouble_receiver
63 DEFINE_L_TYPED_TRANCEIVER(ldouble , long double) // interface i_ldouble_tranceiver
64 DEFINE_C_TYPED_DOUBLE_HANDSHAKE(ldouble , long double)
65 // channel c_ldouble_double_handshake
66 #endif
67
68 void nop(const char*, ...)
69 {
70     /* do nothing */
71 }
72
73 import "c_handshake";
74
75 behavior producer(
76 #ifdef USE_FLOATING_POINT
77     i_float_sender pF,
78     i_double_sender pD,
79     i_ldouble_sender pL,
80 #endif
81     i_char_sender pc,
82     i_short_sender ps,
83     i_int_sender pi,
84     i_llong_sender pl)
85 {
86     void main(void)
87     {
88         char c = DATA;
89         short s = DATA;
90         int i = DATA;
91         long long l = DATA;
92 #ifdef USE_FLOATING_POINT
93         float F = FDATA;
94         double D = FDATA;
95         long double L = FDATA;
96 #endif
97         int n;
98
99         print_time ();
100         if (((char*)&s)[0] == DATA)
101         { printf("Producer: appears to be LITTLE endian\n");
102         }
103         else if (((char*)&s)[1] == DATA)
104         { printf("Producer: appears to be BIG endian\n");
105         }
106         else
107         { printf("Producer: appears to be UNKNOWN endian\n");
108         }
109         for(n=0; n<ITERATIONS; n++)
110         {
111             waitfor(10);
112             print_time ();
113             printf("Producer: sending char c = %d (0x%02x)\n", c, c);

```

```

114     pc.send(c);
115     c++;
116     waitfor(10);
117     print_time();
118     printf("Producer: sending short s = %d (0x%04x)\n", s, s);
119     ps.send(s);
120     s++;
121     waitfor(10);
122     print_time();
123     printf("Producer: sending int i = %d (0x%08x)\n", i, i);
124     pi.send(i);
125     i++;
126     waitfor(10);
127     print_time();
128     printf("Producer: sending llong l = %lld (0x%016llx)\n", l, l);
129     pl.send(l);
130     l++;
131     #ifndef USE_FLOATING_POINT
132         waitfor(10);
133         print_time();
134         printf("Producer: sending float F = %g\n", F);
135         pF.send(F);
136         F += .5;
137         waitfor(10);
138         print_time();
139         printf("Producer: sending double D = %g\n", D);
140         pD.send(D);
141         D += .5;
142         waitfor(10);
143         print_time();
144         printf("Producer: sending ldouble L = %Lg\n", L);
145         pL.send(L);
146         L += .5;
147     #endif
148     }
149     print_time();
150     printf("Producer: done.\n");
151 }
152 };
153
154 behavior consumer(
155     #ifndef USE_FLOATING_POINT
156     i_float_receiver pF,
157     i_double_receiver pD,
158     i_ldouble_receiver pL,
159     #endif
160     i_char_receiver pc,
161     i_short_receiver ps,
162     i_int_receiver pi,
163     i_llong_receiver pl
164     #ifndef EXIT_ON_HANDSHAKE
165     , i_send pdone
166     #endif
167 )
168 {
169     void main(void)
170     {
171         char c;
172         short s = DATA;
173         int i;

```



```

174     long long l;
175 #ifdef USE_FLOATING_POINT
176     float F = FDATA;
177     double D = FDATA;
178     long double L = FDATA;
179 #endif
180     int n;
181
182     print_time ();
183     if (((char*)&s)[0] == DATA)
184     { printf("Consumer: appears to be LITTLE endian\n");
185     }
186     else if (((char*)&s)[1] == DATA)
187     { printf("Consumer: appears to be BIG endian\n");
188     }
189     else
190     { printf("Consumer: appears to be UNKNOWN endian\n");
191     }
192     s = 0;
193     for(n=0; n<ITERATIONS; n++)
194     {
195         pc.receive(&c);
196         print_time ();
197         printf("Consumer: received char c = %d (0x%02x)\n", c, c);
198         ps.receive(&s);
199         print_time ();
200         printf("Consumer: received short s = %d (0x%04x)\n", s, s);
201         pi.receive(&i);
202         print_time ();
203         printf("Consumer: received int i = %d (0x%08x)\n", i, i);
204         pl.receive(&l);
205         print_time ();
206         printf("Consumer: received llong l = %lld (0x%016lx)\n", l, l);
207 #ifdef USE_FLOATING_POINT
208         pF.receive(&F);
209         print_time ();
210         printf("Consumer: received float F = %g\n", F);
211         pD.receive(&D);
212         print_time ();
213         printf("Consumer: received double D = %g\n", D);
214         pL.receive(&L);
215         print_time ();
216         printf("Consumer: received ldouble L = %Lg\n", L);
217 #endif
218     }
219     print_time ();
220     printf("Consumer: done.\n");
221 #ifdef EXIT_ON_HANDSHAKE
222     pdone.send ();
223 #endif
224     }
225 };
226
227 behavior monitor(
228 #ifdef EXIT_ON_HANDSHAKE
229     i_receive pdone
230 #endif
231 )
232 {
233     void main(void)

```

```

234     {
235     #ifdef EXIT_ON_TIME
236         waitfor(EXIT_ON_TIME);
237     #endif
238     #ifdef EXIT_ON_HANDSHAKE
239         pdone.receive();
240     #endif
241         print_time();
242         printf("Monitor: Done, exiting...\n");
243         exit(0);
244     }
245 };
246
247 behavior DUT
248 {
249     c_char_double_handshake cc;
250     c_short_double_handshake cs;
251     c_int_double_handshake ci;
252     c_llong_double_handshake cl;
253     #ifdef USE_FLOATING_POINT
254         c_float_double_handshake cF;
255         c_double_double_handshake cD;
256         c_ldouble_double_handshake cL;
257     #endif
258     #ifdef EXIT_ON_HANDSHAKE
259         c_handshake          cend;
260     #endif
261     producer prod(
262     #ifdef USE_FLOATING_POINT
263         cF, cD, cL,
264     #endif
265         cc, cs, ci, cl);
266     consumer cons(
267     #ifdef USE_FLOATING_POINT
268         cF, cD, cL,
269     #endif
270         cc, cs, ci, cl
271     #ifdef EXIT_ON_HANDSHAKE
272         , cend
273     #endif
274         );
275     monitor mon
276     #ifdef EXIT_ON_HANDSHAKE
277         (cend)
278     #endif
279     ;
280
281     void main(void)
282     {
283         int i;
284         for (i = 0; i < MAXTHREAD; i++)
285             par { prod.main();
286                 cons.main();
287                 mon.main();
288             }
289     }
290 };
291
292 behavior Main
293 {

```

```

294 DUT top;
295
296 int main(void)
297 {
298     print_time();
299     printf("Main: starting...\n");
300     top.main();
301     print_time();
302     printf("Main: done.\n");
303     return 0;
304 }
305 };
306
307 // EOF prodcons.sc

```

A.2 TFMUL Model

Listing 2: TFMUL Model

```

1 // TFMUL.sc: parallel floating-point benchmark
2 // author: Weiwei Chen, Rainer Doemer, Guantao Liu
3 // 02/15/13 GL modified to test HybridThreads library
4 // 11/13/11 RD modified to create more parallel threads
5 // 09/02/11 WC created to test parallel simulators
6
7 #include <stdio.h>
8 #include <stdlib.h>
9 #include <sim.sh>
10
11 // number of multiplications per unit
12 #define MAXLOOP 1000
13
14 // number of threads
15 #ifndef MAXTHREAD
16 #define MAXTHREAD 100000
17 #endif
18
19 // type of floating-point numbers
20 typedef double float_t;
21
22 behavior Fmul
23 {
24     int i = 0;
25     float_t f = 1.2;
26
27     void main()
28     {
29         while(i < MAXLOOP)
30         {
31             f *= 1.1;
32             i++;
33         }
34     }
35 };
36
37 behavior Main
38 {
39     Fmul    fmul0, fmul1, fmul2, fmul3, fmul4,

```

```

40         fmul5, fmul6, fmul7, fmul8, fmul9;
41
42     int main(void)
43     {
44         int i;
45         char *ptr47, *ptr53, *ptr73, *ptr89;
46         printf("Fmul[%d,%d] starting ... \n", MAXTHREAD, MAXLOOP);
47         for(i = 0; i < MAXTHREAD; i++)
48         {
49             par { fmul0; }
50             ptr47 = (char*)malloc(47);
51             par { fmul0; fmul1; fmul2; fmul3; fmul4; fmul5; fmul6; }
52             ptr73 = (char*)malloc(73);
53             free(ptr47);
54             par { fmul0; fmul1; fmul2; fmul3; fmul4; fmul5; fmul6; fmul7; fmul8; }
55             ptr53 = (char*)malloc(53);
56             free(ptr73);
57             par { fmul0; fmul1; fmul2; }
58             ptr89 = (char*)malloc(89);
59             free(ptr53);
60             par { fmul0; fmul1; fmul2; fmul3; fmul4; fmul5; fmul6; fmul7; }
61             ptr73 = (char*)malloc(73);
62             free(ptr89);
63             par { fmul0; fmul1; fmul2; fmul3; }
64             ptr47 = (char*)malloc(47);
65             free(ptr73);
66             par { fmul0; fmul1; }
67             ptr89 = (char*)malloc(89);
68             free(ptr47);
69             par { fmul0; fmul1; fmul2; fmul3; fmul4; }
70             ptr53 = (char*)malloc(53);
71             free(ptr89);
72             par { fmul0; fmul1; fmul2; fmul3; fmul4; fmul5; }
73             ptr73 = (char*)malloc(73);
74             free(ptr53);
75             par { fmul0; fmul1; fmul2; fmul3; fmul4; fmul5; fmul6; fmul7; fmul8; fmul9; }
76             ptr89 = (char*)malloc(89);
77             free(ptr73);
78             par { fmul0; }
79             ptr47 = (char*)malloc(47);
80             free(ptr89);
81             par { fmul0; fmul1; fmul2; fmul3; fmul4; fmul5; fmul6; }
82             ptr73 = (char*)malloc(73);
83             free(ptr47);
84             par { fmul0; fmul1; fmul2; fmul3; fmul4; fmul5; fmul6; fmul7; fmul8; }
85             ptr53 = (char*)malloc(53);
86             free(ptr73);
87             par { fmul0; fmul1; fmul2; }
88             ptr89 = (char*)malloc(89);
89             free(ptr53);
90             par { fmul0; fmul1; fmul2; fmul3; fmul4; fmul5; fmul6; fmul7; }
91             ptr73 = (char*)malloc(73);
92             free(ptr89);
93             par { fmul0; fmul1; fmul2; fmul3; }
94             ptr47 = (char*)malloc(47);
95             free(ptr73);
96             par { fmul0; fmul1; }
97             ptr89 = (char*)malloc(89);
98             free(ptr47);
99             par { fmul0; fmul1; fmul2; fmul3; fmul4; }

```

```

100     ptr53 = (char*)malloc(53);
101     free(ptr89);
102     par { fmul0; fmul1; fmul2; fmul3; fmul4; fmul5; }
103     ptr73 = (char*)malloc(73);
104     free(ptr53);
105     free(ptr73);
106 }
107 printf("Done!\n");
108 return(0);
109 }
110 };
111
112 // EOF TFMUL.sc

```

B Measured Simulation Times for All Benchmarks and Applications

B.1 Simulation Time for Producer-Consumer Model

Table 3: Producer-Consumer Model on mu

Hostname	Usr Time	Sys Time	Elapsed Time	CPU Load	Thread Library
mu	26.78s	0	26.79s	99.00%	QuickThreads
	26.97s	0	26.99s	99.00%	
	26.91s	0	26.92s	99.00%	
	26.85s	0	26.87s	99.00%	
	26.83s	0	26.84s	99.00%	
mu	34.27s	14.04s	48.34s	99.00%	ContextThreads
	34.24s	14.06s	48.32s	99.00%	
	34.54s	14.23s	48.79s	99.00%	
	34.41s	14.1s	48.53s	99.00%	
	34.11s	14.22s	48.35s	99.00%	
mu	84.8s	191.57s	276.46s	99.00%	PosixThreads
	84.49s	189.62s	274.21s	99.00%	
	84.8s	189.48s	274.38s	99.00%	
	84.22s	188.86s	273.18s	99.00%	
	84.16s	191.44s	275.69s	99.00%	

Table 4: Producer-Consumer Model on xi

Hostname	Usr Time	Sys Time	Elapsed Time	CPU Load	Thread Library
xi	22.11s	0	22.17s	99.00%	QuickThreads
	22.08s	0	22.15s	99.00%	
	22.08s	0	22.14s	99.00%	
	22.04s	0	22.11s	99.00%	
	21.8s	0	21.86s	99.00%	
xi	28.44s	10.04s	38.6s	99.00%	ContextThreads
	28.57s	10.05s	38.74s	99.00%	
	28.82	10.28s	39.22s	99.00%	
	28.75s	9.79s	38.65s	99.00%	
	28.1s	10.16s	38.37s	99.00%	
xi	63.86s	233.22s	297.9s	99.00%	PosixThreads
	63.6s	231.25s	295.66s	99.00%	
	65.28s	228.73s	294.82s	99.00%	
	63.05s	229.41s	293.27s	99.00%	
	64.88s	234.53s	300.24s	99.00%	

B.2 Simulation Time for TFMUL Model

Table 5: TFMUL Model on mu

Hostname	Usr Time	Sys Time	Elapsed Time	CPU Load	Thread Library
mu	8.42s	17.97s	26.41s	99.00%	QuickThreads
	8.61s	17.67s	26.29s	99.00%	
	8.28s	17.56s	25.86s	99.00%	
	8.41s	17.73s	26.15s	99.00%	
	8.61s	18.31s	26.94s	99.00%	
mu	11.21s	24.08s	35.31s	99.00%	ContextThreads
	11.1s	24.36s	35.48s	99.00%	
	11.25s	23.68s	34.95s	99.00%	
	11.17s	24.32s	35.51s	99.00%	
	11.73s	24.89s	36.64s	99.00%	
mu	37.6s	170.21s	231.26s	89.00%	PosixThreads
	37.61s	170.56s	231.48s	89.00%	
	38.22s	169.36s	230.67s	89.00%	
	37.38s	169.43s	229.93s	89.00%	
	37.46s	168.77s	229.25s	89.00%	

Table 6: TFMUL Model on xi

Hostname	Usr Time	Sys Time	Elapsed Time	CPU Load	Thread Library
xi	7.38s	11.21s	18.69s	99.00%	QuickThreads
	7.06s	11.34s	18.48s	99.00%	
	7.31s	11.85s	19.27s	99.00%	
	7.11s	12.08s	19.3s	99.00%	
	7.08s	11.37s	18.55s	99.00%	
xi	10.48s	18.51s	29.12s	99.00%	ContextThreads
	10.28s	17.29s	27.7s	99.00%	
	10.1s	16.72s	26.93s	99.00%	
	9.86s	16.36s	26.34s	99.00%	
	9.63s	16.63s	26.39s	99.00%	
xi	39.46s	165.52s	225.7s	90.00%	PosixThreads
	37.84s	163.84s	222.02s	90.00%	
	37.71s	163.21s	221.67s	90.00%	
	37.46s	164.67s	222.81s	90.00%	
	36.42s	162.83s	219.74s	90.00%	