

95

```

1  /*
2  **  NGPT - Next Generation POSIX Threading
3  **  Copyright (c) 2001 IBM Corporation <babt@us.ibm.com>
4  **
5  **  This file is part of NGPT, a non-preemptive thread scheduling
6  **  library which can be found at http://www.ibm.com/developer.
7  **
8  **  This library is free software; you can redistribute it and/or
9  **  modify it under the terms of the GNU Lesser General Public
10 **  License as published by the Free Software Foundation; either
11 **  version 2.1 of the License, or (at your option) any later version.
12 **
13 **  This library is distributed in the hope that it will be useful,
14 **  but WITHOUT ANY WARRANTY; without even the implied warranty of
15 **  MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU
16 **  Lesser General Public License for more details.
17 **
18 **  You should have received a copy of the GNU Lesser General Public
19 **  License along with this library; if not, write to the Free Software
20 **  Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307
21 **  USA.
22 **
23 **  pthreadtypes.h: POSIX Thread ("Pthread") API for NGPT
24 */
25
26                                     /* ``Only those who attempt the absurd
27                                     can achieve the impossible.'`
28                                     -- Unknown                               */
29
30 #if !defined _BITS_TYPES_H && !defined _PTHREAD_H
31 # error "Never include <bits/pthreadtypes.h> directly; use <sys/types.h> instead."
32 #endif
33
34 #ifndef _BITS_PTHREADTYPES_H
35 #define _BITS_PTHREADTYPES_H      1
36
37 #define __need_schedparam
38 #undef sched_param
39 #include <bits/sched.h>
40
41 #ifndef _PTHREAD_QLOCK_DEFINED
42 typedef struct pth_qlock_st pth_qlock_t;
43 typedef struct pth_qlock_st pthread_lock_t;
44 struct pth_qlock_st {
45     int          lock      __attribute__((aligned(16)));
46     int          owner;
47     int          count;
48 };
49 #define _PTHREAD_QLOCK_DEFINED
50 #endif
51
52 struct _pthread_fastlock {
53     long int __status;
54     int __spinlock;
55 };
56
57 #ifndef _PTHREAD_DESCR_DEFINED
58 /* Thread descriptors */
59 typedef struct pthread_st *pthread_descr;
60 #define _PTHREAD_DESCR_DEFINED
61 #endif
62
63
64 #ifdef __USE_XOPEN2K
65 /* POSIX spinlock data type. */
66 typedef volatile int pthread_spinlock_t;
67
68 /* POSIX barrier. */
69 typedef struct {
70     struct _pthread_fastlock __ba_lock; /* Lock to guarantee mutual exclusion */
71     int __ba_required; /* Threads needed for completion */
72     int __ba_present; /* Threads waiting */
73     pthread_descr __ba_waiting; /* Queue of waiting threads */
74 } pthread_barrier_t;
75
76 /* barrier attribute */
77 typedef struct {
78     int __pshared;
79 } pthread_barrierattr_t;
80
81 #endif
82
83 /*
84  * Unprotect namespace, so we can define our own variants now
85  */
86 #undef pthread_t
87 #undef pthread_attr_t
88 #undef pthread_key_t
89 #undef pthread_once_t
90 #undef pthread_mutex_t

```

```
91 #undef pthread_mutexattr_t
92 #undef pthread_cond_t
93 #undef pthread_condattr_t
94 #undef pthread_rwlock_t
95 #undef pthread_rwlockattr_t
96
97 /*
98  * Forward structure definitions.
99  * These are mostly opaque to the application.
100  */
101 struct pthread_st;
102 struct pthread_attr_st;
103 struct pthread_cond_st {
104     void * res0;
105     int res1;
106     int res2;
107 };
108 struct pthread_mutex_st {
109     void * res0;
110     int res1;
111     int res2;
112     int res3;
113     int res4;
114     int res5;
115 };
116 struct pthread_mutexattr_st;
117 struct pthread_rwlock_st;
118
119 /*
120  * Primitive system data type definitions required by P1003.1c
121  */
122 typedef struct pthread_st          *pthread_t;
123 typedef struct pthread_attr_st     *pthread_attr_t;
124 typedef int pthread_key_t;
125 typedef int pthread_once_t;
126 typedef struct pthread_mutexattr_st *pthread_mutexattr_t;
127 typedef struct pthread_mutex_st    pthread_mutex_t;
128 typedef int pthread_condattr_t;
129 typedef struct pthread_cond_st     pthread_cond_t;
130 typedef int pthread_rwlockattr_t;
131 typedef struct pthread_rwlock_st   *pthread_rwlock_t;
132
133 #endif /* bits/pthreadtypes.h */
```

```

1  /*
2  **  NGPT - Next Generation POSIX Threading
3  **  Copyright (c) 2001 IBM Corporation <babt@us.ibm.com>
4  **
5  **  This file is part of NGPT, a non-preemptive thread scheduling
6  **  library which can be found at http://www.ibm.com/developer.
7  **
8  **  This library is free software; you can redistribute it and/or
9  **  modify it under the terms of the GNU Lesser General Public
10 **  License as published by the Free Software Foundation; either
11 **  version 2.1 of the License, or (at your option) any later version.
12 **
13 **  This library is distributed in the hope that it will be useful,
14 **  but WITHOUT ANY WARRANTY; without even the implied warranty of
15 **  MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU
16 **  Lesser General Public License for more details.
17 **
18 **  You should have received a copy of the GNU Lesser General Public
19 **  License along with this library; if not, write to the Free Software
20 **  Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307
21 **  USA.
22 **
23 **  semaphore.h: POSIX Thread ("Pthread") API for NGPT
24 */
25
26                                     /* ``Only those who attempt the absurd
27                                     can achieve the impossible.'`
28                                     -- Unknown          */
29
30 #ifndef _SEMAPHORE_H
31 #define _SEMAPHORE_H    1
32
33 #include <features.h>
34 #include <sys/types.h>
35 #ifdef __USE_XOPEN2K
36 # define __need_timespec
37 # include <time.h>
38 #endif
39
40 #ifndef _PTHREAD_DESCR_DEFINED
41 typedef struct pthread_st *pthread_descr;
42 # define _PTHREAD_DESCR_DEFINED
43 #endif
44
45 typedef struct _sem_lock_st *_sem_lock_t;
46 struct _sem_lock_st {
47     pthread_lock_t    __sem_lock;
48     pthread_mutex_t   __lock;
49     pthread_cond_t    nonzero;
50 };
51
52 /* System specific semaphore definition. */
53 typedef struct {
54     struct _pthread_fastlock    __sem_lock;
55     int                         __sem_value;
56     _sem_lock_t                 __sem_waiting;
57 } sem_t;
58
59 /* Value returned if 'sem_open' failed. */
60 #define SEM_FAILED    ((sem_t *) 0)
61
62 /* Maximum value the semaphore can have. */
63 #define SEM_VALUE_MAX    ((int) ((~0u) >> 1))
64
65 __BEGIN_DECLS
66
67 /* Initialize semaphore object SEM to VALUE. If PSHARED then share it
68 with other processes. */
69 extern int sem_init (sem_t *__sem, int __pshared, unsigned int __value) __THROW;
70
71 /* Free resources associated with semaphore object SEM. */
72 extern int sem_destroy (sem_t *__sem) __THROW;
73
74 /* Open a named semaphore NAME with open flaot OFLAG. */
75 extern sem_t *sem_open (__const char *__name, int __oflag, ...) __THROW;
76
77 /* Close descriptor for named semaphore SEM. */
78 extern int sem_close (sem_t *__sem) __THROW;
79
80 /* Remove named semaphore NAME. */
81 extern int sem_unlink (__const char *__name) __THROW;
82
83 /* Wait for SEM being posted. */
84 extern int sem_wait (sem_t *__sem) __THROW;
85
86 #ifdef __USE_XOPEN2K
87 /* Similar to 'sem_wait' but wait only until ABSTIME. */
88 extern int sem_timedwait (sem_t *__restrict __sem,
89                          __const struct timespec *__restrict __abstime)
90

```

```
91     __THROW;  
92 #endif  
93  
94 /* Test whether SEM is posted. */  
95 extern int sem_trywait (sem_t *__sem) __THROW;  
96  
97 /* Post SEM. */  
98 extern int sem_post (sem_t *__sem) __THROW;  
99  
100 /* Get current value of SEM and store it in *SVAL. */  
101 extern int sem_getvalue (sem_t *__restrict __sem, int *__restrict __sval)  
102     __THROW;  
103  
104 __END_DECLS  
105  
106 #endif /* semaphore.h */
```

```

1  /*
2  *
3  *   Copyright (c) International Business Machines Corp., 2001
4  *
5  *   This program is free software; you can redistribute it and/or modify
6  *   it under the terms of the GNU General Public License as published by
7  *   the Free Software Foundation; either version 2 of the License, or
8  *   (at your option) any later version.
9  *
10 *   This program is distributed in the hope that it will be useful,
11 *   but WITHOUT ANY WARRANTY; without even the implied warranty of
12 *   MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See
13 *   the GNU General Public License for more details.
14 *
15 *   You should have received a copy of the GNU General Public License
16 *   along with this program; if not, write to the Free Software
17 *   Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
18 */
19
20 /*
21 *   FILE      : pth_str01.c
22 *   DESCRIPTION : create a tree of threads
23 *   HISTORY:
24 *       04/09/2001 Paul Larson (plars@us.ibm.com)
25 *       -Ported
26 *
27 */
28
29 #ifndef GLOBAL
30 #include <pthread.h>
31 #else
32 #define _PTHREAD_PRIVATE
33 #include "pthread.h"
34 #endif
35
36
37 #include <stdio.h>
38 #include <unistd.h>
39 #include <stdlib.h>
40 #include <string.h>
41 #include <errno.h>
42 #include <sys/types.h>
43 #include "test_str.h"
44
45 int     depth = 3;
46 int     breadth = 4;
47 int     timeout = 30;           /* minutes */
48 int     cdepth;               /* current depth */
49 int     debug = 0;
50
51 c_info     *child_info;       /* pointer to info array */
52 int     node_count;          /* number of nodes created so far */
53 pthread_mutex_t node_mutex;  /* mutex for the node_count */
54 pthread_cond_t node_condvar; /* condition variable for node_count */
55 pthread_attr_t attr;         /* attributes for created threads */
56
57 int num_nodes(int, int);
58 int synchronize_children(c_info *);
59 int doit(c_info *);
60
61 /*
62 *   parse_args
63 *
64 *   Parse command line arguments. Any errors cause the program to exit
65 *   at this point.
66 */
67 static void parse_args( int argc, char *argv[] )
68 {
69     int     opt, errflag = 0;
70     int     bflag = 0, dflag = 0, tflag = 0;
71
72     while ( (opt = getopt( argc, argv, "b:d:tDh?" )) != EOF ) {
73         switch ( opt ) {
74             case 'b':
75                 if ( bflag )
76                     errflag++;
77                 else {
78                     bflag++;
79                     breadth = atoi( optarg );
80                     if ( breadth <= 0 )
81                         errflag++;
82                 }
83                 break;
84             case 'd':
85                 if ( dflag )
86                     errflag++;
87                 else {
88                     dflag++;
89                     depth = atoi( optarg );
90                     if ( depth <= 0 )

```

```
91         errflag++;
92     }
93     break;
94     case 't':
95         if ( tflag )
96             errflag++;
97         else {
98             tflag++;
99             timeout = atoi( optarg );
100             if ( timeout <= 0 )
101                 errflag++;
102         }
103     break;
104     case 'D':
105         debug = 1;
106     break;
107     case 'h':
108     default:
109         errflag++;
110     break;
111 }
112
113
114 if ( errflag ) {
115     fprintf( stderr, "usage: %s [-b <num>] [-d <num>] [-t <num>] [-D]", argv[0] );
116     fprintf( stderr, " where:\n" );
117     fprintf( stderr, "\t-b <num>\tbreadth of child nodes\n" );
118     fprintf( stderr, "\t-d <num>\tdepth of child nodes\n" );
119     fprintf( stderr, "\t-t <num>\ttimeout for child communication (in minutes)\n" );
120     fprintf( stderr, "\t-D\t\tdebug mode on\n" );
121     exit( 1 );
122 }
123
124 }
125
```

```
126 /*
127  * num_nodes
128  *
129  * Calculate the number of child nodes for a given breadth and depth tree.
130  */
131 int num_nodes( int b, int d )
132 {
133     int n, sum = 1, partial_exp = 1;
134
135     /*
136      * The total number of children = sum ( b ** n )
137      *                               n=0->d
138      * Since b ** 0 == 1, and it's hard to compute that kind of value
139      * in this simplistic loop, we start sum at 1 (above) to compensate
140      * and do the computations from 1->d below.
141     */
142     for ( n = 1; n <= d; n++ ) {
143         partial_exp *= b;
144         sum += partial_exp;
145     }
146
147     return( sum );
148 }
149
```



```

150 /*
151  * synchronize_children
152  *
153  * Register the child with the parent and then wait for all of the children
154  * at the same level to register also. Return when everything is synched up.
155  */
156 int synchronize_children( c_info *parent ) {
157     int          my_index, rc;
158     c_info       *info_p;
159     struct timespec timer;
160
161     if ( debug ) {
162         printf( "trying to lock node_mutex\n" );
163         fflush( stdout );
164     }
165
166     /*
167     * Lock the node_count mutex to we can safely increment it. We
168     * will unlock it when we broadcast that all of our siblings have
169     * been created or when we block waiting for that broadcast.
170     */
171     pthread_mutex_lock( &node_mutex );
172     my_index = node_count++;
173
174     printf( "thread %d started\n", my_index );
175     fflush( stdout );
176
177     /*
178     * Get a pointer into the array of thread structures which will
179     * be "me".
180     */
181     info_p = &child_info[my_index];
182     info_p->index = my_index;
183
184     if ( debug ) {
185         printf( "thread %d info_p=%x\n", my_index, (unsigned int)info_p );
186         fflush( stdout );
187     }
188
189     /*
190     * Register with parent bumping the parent's child_count variable.
191     * Make sure we have exclusive access to that variable before we
192     * do the increment.
193     */
194     if ( debug ) {
195         printf( "thread %d locking child_mutex %x\n", my_index,
196             (unsigned int)&parent->child_mutex );
197         fflush( stdout );
198     }
199     pthread_mutex_lock( &parent->child_mutex );
200     if ( debug ) {
201         printf( "thread %d bumping child_count (currently %d)\n",
202             my_index, parent->child_count );
203         fflush( stdout );
204     }
205     parent->child_ptrs[parent->child_count++] = info_p;
206     if ( debug ) {
207         printf( "thread %d unlocking child_mutex %x\n", my_index,
208             (unsigned int)&parent->child_mutex );
209         fflush( stdout );
210     }
211     pthread_mutex_unlock( &parent->child_mutex );
212
213     if ( debug ) {
214         printf( "thread %d node_count = %d\n", my_index, node_count );
215         printf( "expecting %d nodes\n", num_nodes(breadth, cdepth) );
216         fflush( stdout );
217     }
218
219     /*
220     * Have all the nodes at our level in the tree been created yet?
221     * If so, then send out a broadcast to wake everyone else up (to let
222     * them know they can now create their children (if they are not
223     * leaf nodes)). Otherwise, go to sleep waiting for the broadcast.
224     */
225     if ( node_count == num_nodes(breadth, cdepth) ) {
226
227         /*
228         * Increase the current depth variable, as the tree is now
229         * fully one level taller.
230         */
231         if ( debug ) {
232             printf( "thread %d doing cdepth++(%d)\n", my_index, cdepth );
233             fflush( stdout );
234         }
235         cdepth++;
236
237         if ( debug ) {
238             printf( "thread %d sending child_mutex broadcast\n", my_index );
239             fflush( stdout );

```

```
240     }
241     /*
242     * Since all of our siblings have been created, this level of
243     * the tree is now allowed to continue its work, so wake up the
244     * rest of the siblings.
245     */
246     pthread_cond_broadcast( &node_condvar );
247
248     } else {
249
250     /*
251     * In this case, not all of our siblings at this level of the
252     * tree have been created, so go to sleep and wait for the
253     * broadcast on node_condvar.
254     */
255     if ( debug ) {
256         printf( "thread %d waiting for siblings to register\n",
257             my_index );
258         fflush( stdout );
259     }
260     time( &timer.tv_sec );
261     timer.tv_sec += (unsigned long)timeout * 60;
262     timer.tv_nsec = (unsigned long)0;
263     if ((rc = pthread_cond_timedwait(&node_condvar, &node_mutex,
264         &timer))) {
265         fprintf( stderr, "pthread_cond_timedwait(sync) %d: %s\n",
266             my_index, sys_errlist[rc] );
267         exit( 2 );
268     }
269
270     if ( debug ) {
271         printf( "thread %d is now unblocked\n", my_index );
272         fflush( stdout );
273     }
274
275     }
276
277     /*
278     * Unlock the node_mutex lock, as this thread is finished
279     * initializing.
280     */
281     if ( debug ) {
282         printf( "thread %d unlocking node_mutex\n", my_index );
283         fflush( stdout );
284     }
285     pthread_mutex_unlock( &node_mutex );
286     if ( debug ) {
287         printf( "thread %d unlocked node_mutex\n", my_index );
288         fflush( stdout );
289     }
290
291     if ( debug ) {
292         printf( "synchronize_children returning %d\n", my_index );
293         fflush( stdout );
294     }
295
296     return( my_index );
297 }
298
299
```

```

300  /*
301  * doit
302  *
303  * Do it.
304  */
305  int    doit( c_info *parent ) {
306          int        rc, child, my_index;
307          void        *status;
308          c_info      *info_p;
309          struct timespec timer;
310
311          if ( debug ) {
312              printf( "parent=%#010x\n", (unsigned int)parent );
313              fflush( stdout );
314          }
315
316          if ( parent != NULL ) {
317              /*
318               * Synchronize with our siblings so that all the children at
319               * a given level have been created before we allow those children
320               * to spawn new ones (or do anything else for that matter).
321               */
322              if ( debug ) {
323                  printf( "non-root child calling synchronize_children\n" );
324                  fflush( stdout );
325              }
326              my_index = synchronize_children( parent );
327              if ( debug ) {
328                  printf( "non-root child has been assigned index %d\n",
329                      my_index );
330                  fflush( stdout );
331              }
332          } else {
333              /*
334               * The first thread has no one with which to synchronize, so
335               * set some initial values for things.
336               */
337              if ( debug ) {
338                  printf( "root child\n" );
339                  fflush( stdout );
340              }
341              cdepth = 1;
342              my_index = 0;
343              node_count = 1;
344          }
345
346          /*
347           * Figure out our place in the pthread array.
348           */
349          info_p = &child_info[my_index];
350
351          if ( debug ) {
352              printf( "thread %d getting to heart of doit.\n", my_index );
353              printf( "info_p=%x,cdepth=%d,depth=%d\n", (unsigned int)info_p, cdepth, depth );
354              fflush( stdout );
355          }
356
357          if ( cdepth <= depth ) {
358              /*
359               * Since the tree is not yet complete (it is not yet tall enough),
360               * we need to create another level of children.
361               */
362
363              printf( "thread %d creating kids, cdepth=%d\n", my_index, cdepth );
364              fflush( stdout );
365
366              /*
367               * Create breadth children.
368               */
369              for ( child = 0; child < breadth; child++ ) {
370                  if ( debug ) {
371                      printf( "thread %d making child %d, ptr=%x\n", my_index,
372                          child, (unsigned int)&(info_p->threads[child]) );
373                      fflush( stdout );
374                  }
375                  if ((rc = pthread_create(&(info_p->threads[child]), &attr,
376                      (void *)doit, (void *)info_p)) {
377                      fprintf( stderr, "pthread_create (doit): %s\n",
378                          sys_errlist[rc] );
379                      exit( 3 );
380                  } else {
381                      if ( debug ) {
382                          printf( "pthread_create made thread %x\n",
383                              (unsigned int)&(info_p->threads[child]) );
384                          fflush( stdout );
385                      }
386                  }
387              }
388          }
389

```

```

390     if ( debug ) {
391         printf( "thread %d waits on kids, cdepth=%d\n", my_index,
392             cdepth );
393         fflush( stdout );
394     }
395
396     /*
397     * Wait for our children to finish before we exit ourselves.
398     */
399     for ( child = 0; child < breadth; child++ ) {
400         if ( debug ) {
401             printf( "attempting join on thread %x\n",
402                 (unsigned int)&(info_p->threads[child]) );
403             fflush( stdout );
404         }
405         if ((rc = pthread_join((info_p->threads[child]), &status))) {
406             if ( debug ) {
407                 fprintf( stderr,
408                     "join failed on thread %d, status addr=%x: %s\n",
409                     my_index, (unsigned int)status, sys_errlist[rc] );
410                 fflush( stderr );
411             }
412             exit( 4 );
413         } else {
414             if ( debug ) {
415                 printf( "thread %d joined child %d ok\n", my_index,
416                     child );
417                 fflush( stdout );
418             }
419         }
420     }
421 } else {
422
423     /*
424     * This is the leaf node case. These children don't create
425     * any kids; they just talk amongst themselves.
426     */
427     printf( "thread %d is a leaf node, depth=%d\n", my_index, cdepth );
428     fflush( stdout );
429
430     /*
431     * Talk to siblings (children of the same parent node).
432     */
433     if ( breadth > 1 ) {
434
435         for ( child = 0; child < breadth; child++ ) {
436             /*
437             * Don't talk to yourself.
438             */
439             if ( parent->child_ptrs[child] != info_p ) {
440                 if ( debug ) {
441                     printf( "thread %d locking talk_mutex\n",
442                         my_index );
443                     fflush( stdout );
444                 }
445                 pthread_mutex_lock(
446                     &(parent->child_ptrs[child]->talk_mutex) );
447                 if ( ++parent->child_ptrs[child]->talk_count
448                     == (breadth - 1) ) {
449                     if ( debug ) {
450                         printf( "thread %d talk siblings\n", my_index );
451                         fflush( stdout );
452                     }
453                     if ((rc = pthread_cond_broadcast(
454                         &parent->child_ptrs[child]->talk_condvar))) {
455                         fprintf( stderr, "pthread_cond_broadcast: %s\n",
456                             sys_errlist[rc] );
457                         exit( 5 );
458                     }
459                 }
460                 if ( debug ) {
461                     printf( "thread %d unlocking talk_mutex\n",
462                         my_index );
463                     fflush( stdout );
464                 }
465                 pthread_mutex_unlock(
466                     &(parent->child_ptrs[child]->talk_mutex) );
467             }
468         }
469
470         /*
471         * Wait until the breadth - 1 siblings have contacted us.
472         */
473         if ( debug ) {
474             printf( "thread %d waiting for talk siblings\n",
475                 my_index );
476             fflush( stdout );
477         }
478     }
479

```

```
480     pthread_mutex_lock( &info_p->talk_mutex );
481     if ( info_p->talk_count < (breadth - 1) ) {
482         time( &timer.tv_sec );
483         timer.tv_sec += (unsigned long)timeout * 60;
484         timer.tv_nsec = (unsigned long)0;
485         if ((rc = pthread_cond_timedwait(&info_p->talk_condvar,
486             &info_p->talk_mutex, &timer))) {
487             fprintf( stderr,
488                 "pthread_cond_timedwait (leaf) %d: %s\n",
489                     my_index, sys_errlist[rc] );
490             exit( 6 );
491         }
492     }
493     pthread_mutex_unlock( &info_p->talk_mutex );
494 }
495 }
496 }
497 }
498
499 /*
500  * Our work is done. We're outta here.
501  */
502 printf( "thread %d exiting, depth=%d, status=%d, addr=%x\n", my_index,
503     cdepth, info_p->status, (unsigned int)info_p);
504 fflush( stdout );
505
506 pthread_exit( 0 );
507
508 /*NOTREACHED*/
509 return 0;
510 }
511 }
512 }
```

```
513 /*
514  * main
515  */
516 int main( int argc, char *argv[] ) {
517     int rc, ind, total;
518     pthread_t root_thread;
519
520     parse_args( argc, argv );
521
522     /*
523      * Initialize node mutex.
524      */
525     if ((rc = pthread_mutex_init(&node_mutex, NULL)) ) {
526         fprintf( stderr, "pthread_mutex_init(node_mutex): %s\n",
527             sys_errlist[rc] );
528         exit( 7 );
529     }
530
531     /*
532      * Initialize node condition variable.
533      */
534     if ((rc = pthread_cond_init(&node_condvar, NULL)) ) {
535         fprintf( stderr, "pthread_cond_init(node_condvar): %s\n",
536             sys_errlist[rc] );
537         exit( 8 );
538     }
539
540     /*
541      * Allocate pthread info structure array.
542      */
543     total = num_nodes( breadth, depth );
544     printf( "Allocating %d nodes.\n", total );
545     fflush( stdout );
546     if ( (child_info = (c_info *)malloc( total * sizeof(c_info) ) )
547         == NULL ) {
548         perror( "malloc child_info" );
549         exit( 10 );
550     }
551     bzero( child_info, total * sizeof(c_info) );
552
553     if ( debug ) {
554         printf( "Initializing array for %d children\n", total );
555         fflush( stdout );
556     }
557
558     /*
559      * Allocate array of pthreads descriptors and initialize variables.
560      */
561     for ( ind = 0; ind < total; ind++ ) {
562
563         if ( (child_info[ind].threads =
564             (pthread_t *)malloc( breadth * sizeof(pthread_t) ) )
565             == NULL ) {
566             perror( "malloc threads" );
567             exit( 11 );
568         }
569         bzero( child_info[ind].threads, breadth * sizeof(pthread_t) );
570
571         if ( (child_info[ind].child_ptrs =
572             (c_info **)malloc( breadth * sizeof(c_info * ) ) ) == NULL ) {
573             perror( "malloc child_ptrs" );
574             exit( 12 );
575         }
576         bzero( child_info[ind].child_ptrs,
577             breadth * sizeof(c_info * ) );
578
579         if ((rc = pthread_mutex_init(&child_info[ind].child_mutex,
580             NULL)) ) {
581             fprintf( stderr, "pthread_mutex_init child_mutex: %s\n",
582                 sys_errlist[rc] );
583             exit( 13 );
584         }
585
586         if ((rc = pthread_mutex_init(&child_info[ind].talk_mutex,
587             NULL)) ) {
588             fprintf( stderr, "pthread_mutex_init talk_mutex: %s\n",
589                 sys_errlist[rc] );
590             exit( 14 );
591         }
592
593         if ((rc = pthread_cond_init(&child_info[ind].child_condvar,
594             NULL)) ) {
595             fprintf( stderr,
596                 "pthread_cond_init child_condvar: %s\n",
597                 sys_errlist[rc] );
598             exit( 15 );
599         }
600
601         if ((rc = pthread_cond_init(&child_info[ind].talk_condvar,
602             NULL)) ) {
```

```
603         fprintf( stderr, "pthread_cond_init talk_condvar: %s\n",
604                 sys_errlist[rc] );
605         exit( 16 );
606     }
607
608     if ( debug ) {
609         printf( "Successfully initialized child %d.\n", ind );
610         fflush( stdout );
611     }
612 }
613
614
615 printf( "Creating root thread attributes via pthread_attr_init.\n" );
616 fflush( stdout );
617
618 if ((rc = pthread_attr_init(&attr)) {
619     fprintf( stderr, "pthread_attr_init: %s\n", sys_errlist[rc] );
620     exit( 17 );
621 }
622
623 /*
624  * Make sure that all the thread children we create have the
625  * PTHREAD_CREATE_JOINABLE attribute.  If they don't, then the
626  * pthread_join call will sometimes fail and cause mass confusion.
627  */
628 if ((rc = pthread_attr_setdetachstate(&attr, PTHREAD_CREATE_JOINABLE))
629     ) {
630     fprintf( stderr, "pthread_attr_setdetachstate: %s\n",
631             sys_errlist[rc] );
632     exit( 18 );
633 }
634
635 printf( "Creating root thread via pthread_create.\n" );
636 fflush( stdout );
637
638 if ((rc = pthread_create(&root_thread, &attr, (void *)doit, NULL))) {
639     fprintf( stderr, "pthread_create: %s\n", sys_errlist[rc] );
640     exit( 19 );
641 }
642
643 if ( debug ) {
644     printf( "Doing pthread_join.\n" );
645     fflush( stdout );
646 }
647
648 /*
649  * Wait for the root child to exit.
650  */
651 if (( rc = pthread_join(root_thread, NULL) ) ) {
652     fprintf( stderr, "pthread_join: %s\n", sys_errlist[rc] );
653     exit( 20 );
654 }
655
656 if ( debug ) {
657     printf( "About to pthread_exit.\n" );
658     fflush( stdout );
659 }
660
661 exit( 0 );
662 }
```



```

91     parse_args (argc, argv);
92
93     if(test_limit) {
94         printf ("\n Creating as many threads as possible\n\n");
95     } else {
96         printf ("\n Creating %d threads\n\n", num_threads);
97     }
98     thread (0);
99
100    /*
101     * Program completed successfully...
102     */
103    printf ("\ndone...\n\n");
104    fflush (stdout);
105    exit (0);
106 }
107
108
109 /*-----+
110 /              thread ()
111 / ===== /
112 / Function:  Recursively creates threads while num < num_threads
113 / -----+
114 /
115 +-----*/
116 void *thread (void *parm)
117 {
118     int num = (int) parm;
119     pthread_t  th;
120     pthread_attr_t  attr;
121
122     if (debug) {
123         printf ("\Thread [%d]: new\n", num);
124         fflush (stdout);
125     }
126
127     /*
128      * Create threads while num < num_threads...
129      */
130     if (test_limit || (num < num_threads)) {
131
132         if (pthread_attr_init (&attr))
133             sys_error ("pthread_attr_init failed", __LINE__);
134         if (pthread_attr_setdetachstate (&attr, PTHREAD_CREATE_JOINABLE))
135             sys_error ("pthread_attr_setdetachstate failed", __LINE__);
136         if (pthread_create (&th, &attr, thread, (void *) (num + 1))) {
137             if (test_limit) {
138                 printf ("Testing pthread limit, %d pthreads created.\n", num);
139                 pthread_exit(0);
140             }
141             if (errno == EAGAIN) {
142                 fprintf (stderr, "Thread [%d]: unable to create more threads!\n", num);
143                 return NULL;
144             }
145             else
146                 sys_error ("pthread_create failed", __LINE__);
147         }
148         if (test_limit && ((num % 1000) == 0))
149             printf ("Testing pthread limit: %d pthreads created so far...\n", num);
150         pthread_join (th, (void *) NULL);
151     }
152     if (debug) {
153         printf ("\Thread [%d]: done\n", num);
154         fflush (stdout);
155     }
156
157     pthread_exit(0);
158
159     /*NOTREACHED*/
160     return 0;
161 }
162
163
164 /*-----+
165 /              parse_args ()
166 / ===== /
167 / Function:  Parse the command line arguments & initialize global
168 /           variables.
169 / -----+
170 /
171 +-----*/
172 static void parse_args (int argc, char **argv)
173 {
174     int  i;
175     int  errflag = 0;
176     char *program_name = *argv;
177
178     while ((i = getopt(argc, argv, "dl:?:")) != EOF) {
179         switch (i) {
180             case 'd':
181                 /* debug option */

```

```
181         debug++;
182         break;
183     case 'l':          /* test pthread limit */
184         test_limit++;
185         break;
186     case 'n':          /* number of threads */
187         num_threads = atoi (optarg);
188         break;
189     case '?':
190         errflag++;
191         break;
192     }
193 }
194
195 /* If any errors exit program */
196 if (errflag) {
197     fprintf (stderr, USAGE, program_name);
198     exit (2);
199 }
200 }
201
202 /*-----+
203 /          sys_error ()
204 /-----+
205 / Function:  Creates system error message and calls error ()
206 /-----+
207 /-----+
208 /-----+
209 +-----*/
210 static void sys_error (const char *msg, int line)
211 {
212     char syserr_msg [256];
213
214     sprintf (syserr_msg, "%s:%s\n", msg, strerror (errno));
215     error (syserr_msg, line);
216 }
217
218
219 /*-----+
220 /          error ()
221 /-----+
222 /-----+
223 / Function:  Prints out message and exits...
224 /-----+
225 +-----*/
226 static void error (const char *msg, int line)
227 {
228     fprintf (stderr, "ERROR [line: %d] %s\n", line, msg);
229     exit (-1);
230 }
```

```

1  /*
2  *
3  *   Copyright (c) International Business Machines Corp., 2001
4  *
5  *   This program is free software; you can redistribute it and/or modify
6  *   it under the terms of the GNU General Public License as published by
7  *   the Free Software Foundation; either version 2 of the License, or
8  *   (at your option) any later version.
9  *
10 *   This program is distributed in the hope that it will be useful,
11 *   but WITHOUT ANY WARRANTY; without even the implied warranty of
12 *   MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See
13 *   the GNU General Public License for more details.
14 *
15 *   You should have received a copy of the GNU General Public License
16 *   along with this program; if not, write to the Free Software
17 *   Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
18 */
19
20 /*
21 *   FILE      : pth_str01.c
22 *   DESCRIPTION : create a tree of threads does calculations, and
23 *                 returns result to parent
24 *   HISTORY:
25 *       05/16/2001 Paul Larson (plars@us.ibm.com)
26 *       -Ported
27 *
28 */
29
30 #ifndef GLOBAL
31 #include <pthread.h>
32 #else
33 #define _PTHREAD_PRIVATE
34 #include "pthread.h"
35 #endif
36
37 #include <stdio.h>
38 #include <unistd.h>
39 #include <stdlib.h>
40 #include <string.h>
41 #include <errno.h>
42 #include <fcntl.h>
43 #include <sys/types.h>
44 #include <sys/stat.h>
45 #include <sys/mman.h>
46
47 #define MAXTHREADS 50000
48
49 /* Type definition */
50 struct kid_info {
51     long          sum;          /* sum of childrens indexes plus our own */
52     int           index;       /* our index into the array */
53     int           status;      /* return status of this thread */
54     int           child_count; /* Count of children created */
55     int           talk_count;  /* Count of siblings that we talked to */
56     pthread_t     *threads;    /* dynamic array of thread id of kids */
57     pthread_mutex_t talk_mutex; /* mutex for the talk_count */
58     pthread_mutex_t child_mutex; /* mutex for the child_count */
59     pthread_cond_t talk_condvar; /* condition variable for talk_count */
60     pthread_cond_t child_condvar; /* condition variable for child_count */
61     struct kid_info **child_ptrs; /* dynamic array of ptrs to kids */
62 };
63 typedef struct kid_info c_info;
64
65 /* Global variables */
66 int     depth = 3;
67 int     breadth = 4;
68 int     timeout = 30;          /* minutes */
69 int     cdepth;               /* current depth */
70 int     debug = 0;
71 int     shared_fd = 0;
72
73
74 c_info     *child_info;      /* pointer to info array */
75 int     node_count;         /* number of nodes created so far */
76 pthread_mutex_t node_mutex; /* mutex for the node_count */
77 pthread_cond_t node_condvar; /* condition variable for node_count */
78 pthread_attr_t attr;        /* attributes for created threads */
79
80 int     use_shared = 0;
81 typedef struct shared_area shared_area_t;
82 struct shared_area {
83     int     initialized;
84     int     locked;
85     pthread_mutex_t mutex;
86 };
87
88 int num_nodes(int, int);
89 int synchronize_children(c_info *) ;
90 void *doit(void *);

```

```

91
92 /*-----*
93  * parse_args
94  *
95  * Parse command line arguments. Any errors cause the program to exit
96  * at this point.
97  *-----*/
98 static void parse_args( int argc, char *argv[] )
99 {
100     int          opt, errflag = 0;
101     int          bflag = 0, dflag = 0, tflag = 0;
102
103     while ( (opt = getopt( argc, argv, "b:d:tDs?" )) != EOF ) {
104         switch ( opt ) {
105             case 's':
106                 use_shared = 1;
107                 break;
108             case 'b':
109                 if ( bflag )
110                     errflag++;
111                 else {
112                     bflag++;
113                     breadth = atoi( optarg );
114                     if ( breadth <= 0 )
115                         errflag++;
116                 }
117                 break;
118             case 'd':
119                 if ( dflag )
120                     errflag++;
121                 else {
122                     dflag++;
123                     depth = atoi( optarg );
124                     if ( depth <= 0 )
125                         errflag++;
126                 }
127                 break;
128             case 't':
129                 if ( tflag )
130                     errflag++;
131                 else {
132                     tflag++;
133                     timeout = atoi( optarg );
134                     if ( timeout <= 0 )
135                         errflag++;
136                 }
137                 break;
138             case 'D':
139                 debug = 1;
140                 break;
141             case '?':
142             default:
143                 errflag++;
144                 break;
145         }
146     }
147
148     if ( errflag ) {
149         fprintf( stderr, "usage: %s [-b <num>] [-d <num>] [-t <num>] [-s] [-D]", argv[0] );
150         fprintf( stderr, "where:\n" );
151         fprintf( stderr, "\t-b <num>\tbreadth of child nodes\n" );
152         fprintf( stderr, "\t-d <num>\tdepth of child nodes\n" );
153         fprintf( stderr, "\t-t <num>\ttimeout for child communication (in minutes)\n" );
154         fprintf( stderr, "\t-s \t\twrap with shared mutex\n" );
155         fprintf( stderr, "\t-D\t\tdebug mode\n" );
156         exit( 1 );
157     }
158 }
159
160
161
162 /*-----*
163  * num_nodes
164  *
165  * Caculate the number of child nodes for a given breadth and depth tree.
166  *-----*/
167 int num_nodes( int b, int d )
168 {
169     int          n, sum = 1, partial_exp = 1;
170
171     /*
172     * The total number of children = sum ( b ** n )
173     *                               n=0->d
174     * Since b ** 0 == 1, and it's hard to compute that kind of value
175     * in this simplistic loop, we start sum at 1 (above) to compensate
176     * and do the computations from 1->d below.
177     */
178     for ( n = 1; n <= d; n++ ) {
179         partial_exp *= b;
180         sum += partial_exp;

```

```

181     }
182
183     return( sum );
184 }
185
186
187 /*-----*
188  * synchronize_children
189  *
190  * Register the child with the parent and then wait for all of the children
191  * at the same level to register also. Return when everything is synched up.
192  *-----*/
193 int synchronize_children( c_info *parent )
194 {
195     int             my_index, rc;
196     c_info          *info_p;
197     struct timespec timer;
198
199     if ( debug ) {
200         printf( "trying to lock node_mutex\n" );
201         fflush( stdout );
202     }
203
204     /* Lock the node_count mutex to we can safely increment it. We
205      * will unlock it when we broadcast that all of our siblings have
206      * been created or when we block waiting for that broadcast. */
207     pthread_mutex_lock( &node_mutex );
208     my_index = node_count++;
209
210     printf( "thread %d started\n", my_index );
211     fflush( stdout );
212
213     /* Get a pointer into the array of thread structures which will be "me". */
214     info_p = &child_info[my_index];
215     info_p->index = my_index;
216     info_p->sum = (long) my_index;
217
218     if ( debug ) {
219         printf( "thread %d info_p=%x\n", my_index, (unsigned int)info_p );
220         fflush( stdout );
221     }
222
223     /* Register with parent bumping the parent's child_count variable.
224      * Make sure we have exclusive access to that variable before we
225      * do the increment. */
226     if ( debug ) {
227         printf( "thread %d locking child_mutex %x\n", my_index, (unsigned int)&parent->child_mutex );
228         fflush( stdout );
229     }
230     pthread_mutex_lock( &parent->child_mutex );
231     if ( debug ) {
232         printf( "thread %d bumping child_count (currently %d)\n",
233             my_index, parent->child_count );
234         fflush( stdout );
235     }
236     parent->child_ptrs[parent->child_count++] = info_p;
237     if ( debug ) {
238         printf( "thread %d unlocking child_mutex %x\n", my_index,
239             (unsigned int)&parent->child_mutex );
240         fflush( stdout );
241     }
242     pthread_mutex_unlock( &parent->child_mutex );
243
244     if ( debug ) {
245         printf( "thread %d node_count = %d\n", my_index, node_count );
246         printf( "expecting %d nodes\n", num_nodes(breadth, cdepth) );
247         fflush( stdout );
248     }
249
250     /* Have all the nodes at our level in the tree been created yet?
251      * If so, then send out a broadcast to wake everyone else up (to let
252      * them know they can now create their children (if they are not
253      * leaf nodes)). Otherwise, go to sleep waiting for the broadcast. */
254     if ( node_count == num_nodes(breadth, cdepth) ) {
255
256         /* Increase the current depth variable, as the tree is now
257          * fully one level taller. */
258         if ( debug ) {
259             printf( "thread %d doing cdepth++ (%d)\n", my_index, cdepth );
260             fflush( stdout );
261         }
262         cdepth++;
263
264         if ( debug ) {
265             printf( "thread %d sending child_mutex broadcast\n", my_index );
266             fflush( stdout );
267         }
268
269         /* Since all of our siblings have been created, this level of
270          * the tree is now allowed to continue its work, so wake up the

```

```

271         * rest of the siblings. */
272         pthread_cond_broadcast( &node_condvar );
273
274     } else {
275
276         /* In this case, not all of our siblings at this level of the
277          * tree have been created, so go to sleep and wait for the
278          * broadcast on node_condvar. */
279         if ( debug ) {
280             printf( "thread %d waiting for siblings to register\n",
281                 my_index );
282             fflush( stdout );
283         }
284         time( &timer.tv_sec );
285         timer.tv_sec += (unsigned long)timeout * 60;
286         timer.tv_nsec = (unsigned long)0;
287         if ((rc = pthread_cond_timedwait(&node_condvar, &node_mutex,
288             &timer))) {
289             fprintf( stderr, "pthread_cond_timedwait(sync) %d: %s\n",
290                 my_index, sys_errlist[rc] );
291             exit( 2 );
292         }
293
294         if ( debug ) {
295             printf( "thread %d is now unblocked\n", my_index );
296             fflush( stdout );
297         }
298     }
299
300     /* Unlock the node_mutex lock, as this thread is finished initializing. */
301     if ( debug ) {
302         printf( "thread %d unlocking node_mutex\n", my_index );
303         fflush( stdout );
304     }
305     pthread_mutex_unlock( &node_mutex );
306     if ( debug ) {
307         printf( "thread %d unlocked node_mutex\n", my_index );
308         fflush( stdout );
309     }
310
311     if ( debug ) {
312         printf( "synchronize_children returning %d\n", my_index );
313         fflush( stdout );
314     }
315
316     return( my_index );
317 }
318
319
320
321 /*-----*
322  * doit
323  *
324  * Do it.
325  *-----*/
326 void *doit( void *param )
327 {
328     c_info *parent = (c_info *) param;
329     int     rc, child, my_index;
330     void    *status;
331     c_info  *info_p;
332     struct timespec timer;
333
334     if ( debug ) {
335         printf( "parent=%#010x\n", (unsigned int)parent );
336         fflush( stdout );
337     }
338
339     if ( parent != NULL ) {
340         /* Synchronize with our siblings so that all the children at
341          * a given level have been created before we allow those children
342          * to spawn new ones (or do anything else for that matter). */
343         if ( debug ) {
344             printf( "non-root child calling synchronize_children\n" );
345             fflush( stdout );
346         }
347         my_index = synchronize_children( parent );
348         if ( debug ) {
349             printf( "non-root child has been assigned index %d\n",
350                 my_index );
351             fflush( stdout );
352         }
353     } else {
354         /* The first thread has no one with which to synchronize, so
355          * set some initial values for things. */
356         if ( debug ) {
357             printf( "root child\n" );
358             fflush( stdout );
359         }
360         cdepth = 1;

```

```

361     my_index = 0;
362     node_count = 1;
363 }
364
365 /* Figure out our place in the pthread array. */
366 info_p = &child_info[my_index];
367
368 if ( debug ) {
369     printf( "thread %d getting to heart of doit.\n", my_index );
370     printf( "info_p=%x,cdepth=%d,depth=%d\n", (unsigned int)info_p, cdepth, depth );
371     fflush( stdout );
372 }
373
374 if ( cdepth <= depth )
375 {
376     /* Since the tree is not yet complete (it is not yet tall enough),
377      * we need to create another level of children. */
378
379     printf( "thread %d creating kids,cdepth=%d\n", my_index, cdepth );
380     fflush( stdout );
381
382     /* Create breadth children. */
383     for ( child = 0; child < breadth; child++ ) {
384         if ( debug ) {
385             printf( "thread %d making child %d,ptr=%x\n", my_index,
386                 child, (unsigned int)&(info_p->threads[child]) );
387             fflush( stdout );
388         }
389         if ( (rc = pthread_create(&(info_p->threads[child]), &attr, doit, (void *) info_p)) ) {
390             fprintf( stderr, "pthread_create(doit): %s\n",
391                 sys_errlist[rc] );
392             exit( 3 );
393         } else {
394             if ( debug ) {
395                 printf( "pthread_create made thread %x\n",
396                     (unsigned int)&(info_p->threads[child]) );
397                 fflush( stdout );
398             }
399         }
400     }
401
402     if ( debug ) {
403         printf( "thread %d waits on kids,cdepth=%d\n", my_index,
404             cdepth );
405         fflush( stdout );
406     }
407
408     /* Wait for our children to finish before we exit ourselves. */
409     for ( child = 0; child < breadth; child++ ) {
410         if ( debug ) {
411             printf( "attempting join on thread %x\n",
412                 (unsigned int)&(info_p->threads[child]) );
413             fflush( stdout );
414         }
415         if ( (rc = pthread_join((info_p->threads[child]), &status)) ) {
416             if ( debug ) {
417                 fprintf( stderr,
418                     "join failed on thread %d,status addr=%x: %s\n",
419                     my_index, (unsigned int)status, sys_errlist[rc] );
420                 fflush( stderr );
421             }
422             exit( 4 );
423         } else {
424             if ( debug ) {
425                 printf( "thread %d joined child %d ok\n", my_index, child );
426                 fflush( stdout );
427             }
428
429             /* Add all childrens indexes to your index value */
430             info_p->sum += info_p->child_ptrs[child]->sum;
431             printf( "thread %d adding child thread %d to sum = %d\n",
432                 my_index, info_p->child_ptrs[child]->index, (long int)info_p->sum);
433         }
434     }
435 } else {
436
437     /* This is the leaf node case. These children don't create
438      * any kids; they just talk amongst themselves. */
439     printf( "thread %d is a leaf node,depth=%d\n", my_index, cdepth );
440     fflush( stdout );
441
442     /* Talk to siblings (children of the same parent node). */
443     if ( breadth > 1 ) {
444
445         for ( child = 0; child < breadth; child++ ) {
446             /* Don't talk to yourself. */
447             if ( parent->child_ptrs[child] != info_p ) {
448                 if ( debug ) {
449                     printf( "thread %d locking talk_mutex\n",

```

```

451         my_index );
452         fflush( stdout );
453     }
454     pthread_mutex_lock(
455         &(parent->child_ptrs[child]->talk_mutex) );
456     if ( ++parent->child_ptrs[child]->talk_count == (breadth - 1) ) {
457         if ( debug ) {
458             printf( "thread %d talk siblings\n", my_index );
459             fflush( stdout );
460         }
461         if ((rc = pthread_cond_broadcast(
462             &parent->child_ptrs[child]->talk_condvar)) {
463             fprintf( stderr, "pthread_cond_broadcast: %s\n",
464                 sys_errlist[rc] );
465             exit( 5 );
466         }
467     }
468     if ( debug ) {
469         printf( "thread %d unlocking talk_mutex\n",
470             my_index );
471         fflush( stdout );
472     }
473     pthread_mutex_unlock(
474         &(parent->child_ptrs[child]->talk_mutex) );
475 }
476 }
477
478 /* Wait until the breadth - 1 siblings have contacted us. */
479 if ( debug ) {
480     printf( "thread %d waiting for talk siblings\n", my_index );
481     fflush( stdout );
482 }
483
484 pthread_mutex_lock( &info_p->talk_mutex );
485 if ( info_p->talk_count < (breadth - 1) ) {
486     time( &timer.tv_sec );
487     timer.tv_sec += (unsigned long)timeout * 60;
488     timer.tv_nsec = (unsigned long)0;
489     if ((rc = pthread_cond_timedwait(&info_p->talk_condvar,
490         &info_p->talk_mutex, &timer)) {
491         fprintf( stderr,
492             "pthread_cond_timedwait (leaf %d: %s\n",
493                 my_index, sys_errlist[rc] );
494         exit( 6 );
495     }
496 }
497 pthread_mutex_unlock( &info_p->talk_mutex );
498 }
499 }
500 }
501 }
502
503 /* Our work is done. We're outta here. */
504 printf( "thread %d exiting, depth=%d, status=%d, addr=%x\n", my_index,
505     cdepth, info_p->status, (unsigned int)info_p);
506 fflush( stdout );
507
508 pthread_exit( 0 );
509
510 /*NOTREACHED*/
511 return 0;
512 }
513
514
515
516 /*-----*
517 * main
518 *-----*/
519 int main( int argc, char *argv[] )
520 {
521     int rc, ind, total, i_initialized = 0;
522     pthread_mutexattr_t myattr;
523     pthread_t root_thread;
524     shared_area_t *sa = NULL;
525
526     parse_args( argc, argv );
527
528     if (use_shared == 1) {
529         if (shared_fd == 0) {
530             shared_fd = open("test_str03m", O_RDWR, (S_IRWXU | S_IRWXG | S_IRWXO));
531             if (shared_fd < 0) {
532                 fprintf(stderr, "unable to open shared memory\n");
533                 exit(1);
534             }
535         }
536         sa = (shared_area_t *)mmap(NULL, (sizeof(struct shared_area)),
537             PROT_READ | PROT_WRITE | PROT_EXEC, MAP_SHARED,
538             shared_fd, 0);
539         if (sa == MAP_FAILED) {
540             fprintf(stderr, "unable to map shared memory\n");

```



```

541         exit(1);
542     }
543 }
544
545
546 /* Initialize shared mutex... */
547 if (sa->initialized != 1) {
548     fprintf(stderr, "shared memory was not initialized.\n");
549     pthread_mutexattr_init(&myattr);
550     pthread_mutexattr_setpshared(&myattr, PTHREAD_PROCESS_SHARED);
551     if ((rc = pthread_mutex_init(&(sa->mutex), &myattr)) {
552         fprintf( stderr, "pthread_mutex_init(test_mutex): %s\n",
553                 sys_errlist[rc] );
554         exit( 7 );
555     }
556     fprintf(stderr, "mutex was initialized.\n");
557     sa->initialized = 1;
558     i_initialized = 1;
559     msync(sa, sizeof(struct shared_area), MS_INVALIDATE | MS_SYNC);
560 } else
561     fprintf(stderr, "shared memory was initialized.\n");
562
563 if (i_initialized == 1) {
564     fprintf(stderr, "I did initialize, spinning...\n");
565     fprintf(stderr, "waiting for partner to lock.\n");
566     while (sa->locked != 1) {
567         sleep(1);
568         continue;
569     }
570     fprintf(stderr, "partner is locked.\n");
571 }
572
573 if ((rc = pthread_mutex_lock(&(sa->mutex))) {
574     fprintf(stderr, "unable to grab test_mutex, rc = %d\n", rc);
575     exit( 7 );
576 }
577 fprintf(stderr, "I am locked.\n");
578 sa->locked = 1;
579 msync(sa, sizeof(struct shared_area), MS_INVALIDATE | MS_SYNC);
580
581 }
582
583
584 /* Initialize node mutex. */
585 if ((rc = pthread_mutex_init(&node_mutex, NULL)) {
586     fprintf( stderr, "pthread_mutex_init(node_mutex): %s\n",
587             sys_errlist[rc] );
588     exit( 7 );
589 }
590
591 /* Initialize node condition variable. */
592 if ((rc = pthread_cond_init(&node_condvar, NULL)) {
593     fprintf( stderr, "pthread_cond_init(node_condvar): %s\n",
594             sys_errlist[rc] );
595     exit( 8 );
596 }
597
598 /* Allocate pthread info structure array. */
599 if ( (total = num_nodes( breadth, depth )) > MAXTHREADS ) {
600     fprintf( stderr, "Can't create %d threads; max is %d\n",
601             total, MAXTHREADS );
602     exit( 9 );
603 }
604 printf( "Allocating %d nodes.\n", total );
605 fflush( stdout );
606 if ( (child_info = (c_info *)malloc( total * sizeof(c_info) )) == NULL ) {
607     perror( "malloc child_info" );
608     exit( 10 );
609 }
610 bzero( child_info, total * sizeof(c_info) );
611
612 if ( debug ) {
613     printf( "Initializing array for %d children\n", total );
614     fflush( stdout );
615 }
616
617 /* Allocate array of pthreads descriptors and initialize variables. */
618 for ( ind = 0; ind < total; ind++ ) {
619
620     if ( (child_info[ind].threads =
621          (pthread_t *)malloc( breadth * sizeof(pthread_t) )) == NULL ) {
622         perror( "malloc threads" );
623         exit( 11 );
624     }
625     bzero( child_info[ind].threads, breadth * sizeof(pthread_t) );
626
627     if ( (child_info[ind].child_ptrs =
628          (c_info **)malloc( breadth * sizeof(c_info * ) )) == NULL ) {
629         perror( "malloc child_ptrs" );
630         exit( 12 );

```

```

631     }
632     bzero( child_info[ind].child_ptrs,
633           breadth * sizeof(c_info * ) );
634
635     if ((rc = pthread_mutex_init(&child_info[ind].child_mutex, NULL))) {
636         fprintf( stderr, "pthread_mutex_init child_mutex: %s\n",
637                sys_errlist[rc] );
638         exit( 13 );
639     }
640
641     if ((rc = pthread_mutex_init(&child_info[ind].talk_mutex, NULL))) {
642         fprintf( stderr, "pthread_mutex_init talk_mutex: %s\n",
643                sys_errlist[rc] );
644         exit( 14 );
645     }
646
647     if ((rc = pthread_cond_init(&child_info[ind].child_condvar, NULL))) {
648         fprintf( stderr,
649                "pthread_cond_init child_condvar: %s\n",
650                sys_errlist[rc] );
651         exit( 15 );
652     }
653
654     if ((rc = pthread_cond_init(&child_info[ind].talk_condvar, NULL))) {
655         fprintf( stderr, "pthread_cond_init talk_condvar: %s\n",
656                sys_errlist[rc] );
657         exit( 16 );
658     }
659
660     if ( debug ) {
661         printf( "Successfully initialized child %d\n", ind );
662         fflush( stdout );
663     }
664 }
665
666 printf( "Creating root thread attributes via pthread_attr_init.\n" );
667 fflush( stdout );
668
669 if ((rc = pthread_attr_init(&attr)) {
670     fprintf( stderr, "pthread_attr_init: %s\n", sys_errlist[rc] );
671     exit( 17 );
672 }
673
674 /* Make sure that all the thread children we create have the
675  * PTHREAD_CREATE_JOINABLE attribute. If they don't, then the
676  * pthread_join call will sometimes fail and cause mass confusion. */
677 if ((rc = pthread_attr_setdetachstate(&attr, PTHREAD_CREATE_JOINABLE)) {
678     fprintf( stderr, "pthread_attr_setdetachstate: %s\n",
679            sys_errlist[rc] );
680     exit( 18 );
681 }
682
683 printf( "Creating root thread via pthread_create.\n" );
684 fflush( stdout );
685
686 if ((rc = pthread_create(&root_thread, &attr, doit, NULL)) {
687     fprintf( stderr, "pthread_create: %s\n", sys_errlist[rc] );
688     exit( 19 );
689 }
690
691 if ( debug ) {
692     printf( "Doing pthread_join.\n" );
693     fflush( stdout );
694 }
695
696 /* Wait for the root child to exit. */
697 if ((rc = pthread_join(root_thread, NULL)) {
698     fprintf( stderr, "pthread_join: %s\n", sys_errlist[rc] );
699     exit( 20 );
700 }
701
702 if ( debug ) {
703     printf( "About to pthread_exit.\n" );
704     fflush( stdout );
705 }
706
707 printf( "\n\nThe sum of tree (breadth %d, depth %d) is %ld\n",
708        breadth, depth, child_info[0].sum);
709
710 if (use_shared == 1) {
711     pthread_mutex_unlock(&(sa->mutex));
712     if (i_initialized == 1) {
713         pthread_mutex_destroy(&(sa->mutex));
714         sa->locked = 0;
715         sa->initialized = 0;
716     }
717 }
718 exit( 0 );
719 }
720 }

```

1	./NGPT/ngpt-2.0.1/pthreadtypes.h.....	Pages	1- 2	134 lines
2	./NGPT/ngpt-2.0.1/semaphore.h.....	Pages	3- 4	107 lines
3	./NGPT/ngpt-2.0.1/test_str01.c.....	Pages	5- 14	663 lines
4	./NGPT/ngpt-2.0.1/test_str02.c.....	Pages	15- 17	231 lines
5	./NGPT/ngpt-2.0.1/test_str03.c.....	Pages	18- 25	721 lines

End of Table of Contents