

SMPL Programming Language Proposal

Andrei Papancea, Ajay Siva Santosh, Devashi Tandon
{alp2200, ac3647, dt2450} @columbia.edu

September 16, 2013

Motivation

Up until 2006, computer scientists benefited from what is referred to as the “free lunch”¹. Programmers did not have to focus too much on writing very efficient code, since it was known that every two years processor speeds will increase and hence help their programs run faster. According to Moore’s Law² the number of transistors on integrated circuits doubles approximately every two years. This law still holds today, but it is not practical to increase processor speed anymore because of financial and physical constraints such as the high cost of dealing with excessive heat dissipation. Subsequently, in a multicore world, the computer scientists all over are looking at parallel programming languages as the solution.

Two issues with parallel programming are that it is not trivial given the current toolsets and that not every problem can be parallelized. Furthermore, most popular programming languages such as C, C++, Java, or Python either do not support parallel programming or they have a very complex implementation of it, since they were not built with that idea in mind. Due to the non-trivial syntax of most parallel programming -capable languages, the concept of parallel programming is rarely taught at the undergraduate level — and it is crucial to teach this concept at the undergraduate level in order to stir interest in the field and push for advancement. Hence, there is a need for programming languages with simpler syntax, that can better illustrate parallel programming concepts and that makes them more suitable for undergraduate teaching³.

Language Description

SMPL is an approach to parallel programming that is based on a subset of C. Its major strength is that it has clearer and simpler syntax for doing parallel programming than programs written using OpenMP or POSIX threads.

A SMPL program consists of a sequence of function, variable, and array definitions, similar to C.

¹ Sutter, Herb. "The free lunch is over: A fundamental turn toward concurrency in software." *Dr. Dobbs's Journal* 30.3 (2005): 202-210.

² Schaller, Robert R. "Moore's law: past, present and future." *Spectrum, IEEE* 34.6 (1997): 52-59.

³ Graf, Michael et al. "Selecting and using a parallel programming language." *Proceeding of the 44th ACM technical symposium on Computer science education* 6 Mar. 2013: 735-735.

Syntax Description

The syntax of a SMPL program is like in C. There are 4 parallel constructs in SMPL:

1. `pfor`
 - a. **Syntax:** `pfor(num_threads; initializer; stop condition; updater) { statements }`
 - b. **Definition:** breaks the work of the for loop into the specified number of threads
 - c. **Features:** includes implicit thread exit and join
2. `spawn`
 - a. **Syntax:** `spawn { statement }`
 - b. **Definition:** spawns the specified statement into a separate thread, causing it to run asynchronously
3. `lock`
 - a. **Syntax:** `lock { statement }`
 - b. **Definition:** the C-equivalent of a *mutex* lock, which basically protects the operations inside it from thread concurrency
4. `barrier`
 - a. **Syntax:** `barrier;`
 - b. **Definition:** prevents code below the *barrier* from executing until the currently *spawn*-ed threads have finished

Syntax Examples

Example 1: *Sample Syntax to calculate the sum of first n numbers and checking if there are greater than max without using parallelization.*

```
1 void main(){
2
3     /* This is a multi line
4        comment. Nesting is not
5        supported at this
6        point. */
7
8     int i=0;
9     int sum=0;
10    int n;
11    int max;
12
13    printf("Enter the value of n and max \n");
14    scanf("%d,%d", &n, &max);
15
16    for(i=0;i<n;i++){
17        sum += i;
18    }
19
20    if(sum>max){
21        printf("The sum is greater than max");
22    } else {
23        printf("The sum is not greater than max");
24    }
25
26 }
27
```

Example 2: *Sample Syntax to calculate the sum of first n numbers and checking if there are greater than max using functions without using parallelization.*

```
1 void main() {
2
3     int i=0;
4     int n;
5     int max;
6
7     printf("Enter the value of n and max \n");
8     scanf("%d,%d", &n, &max);
9
10    int sum=F1(n);
11
12    if(sum>max){
13        printf("The sum is greater than max");
14    } else {
15        printf("The sum is not greater than max");
16    }
17
18 }
19
20 int F1(int n){
21
22     int i=0;
23     int sum=0;
24     for(i=0;i<n;i++){
25         sum+=i;
26     }
27     return sum;
28
29 }
30
```

Parallelism in SMPL

The program below computes the sum of the first n positive integers using t threads, where n and t are inputted by the user:

```
1 void main(){
2     int i=0;
3     int sum=0;
4     int n;
5     int t;
6
7     printf("Enter the number of threads t: ");
8     scanf("%d", &t);
9
10    printf("Enter the upper bound n: ");
11    scanf("%d", &n);
12
13    pfor(t; i=0; i<n; i++)
14        lock {
15            sum += i;
16        }
17    }
18
19    printf("The sum of the first %d numbers is %d.\n",
20          limit, sum);
21 }
```

The C-equivalent of the SMPL program above:

```
1 #include <stdio.h>
2 #include <stdlib.h>
3 #include <pthread.h>
4
5 void *thread_sum(void *);
6 pthread_mutex_t m=PTHREAD_MUTEX_INITIALIZER;
7 int sum=0;
8
9 void main(){
10     int i,j;
11     int limit = 1000000;
12     int num_threads;
13
14     printf("Enter number of threads: ");
15     scanf("%d",&num_threads);
16
17     pthread_t threads[num_threads];
18     int uppers[num_threads];
```

```

19  int lowers[num_threads];
20  int args[num_threads][2];
21
22  printf("Enter upper bound: ");
23  scanf("%d",&limit);
24
25  for(i=0;i<num_threads;i++){
26      uppers[i] = (limit/num_threads)*(i+1);
27      lowers[i] = (limit/num_threads)*i;
28      args[i][0] = lowers[i];
29      args[i][1] = uppers[i];
30  }
31
32  for(i=0;i<num_threads;i++){
33      pthread_create(&threads[i],NULL,
34                  thread_sum,(void *)&args[i]);
35  }
36
37  for(i=0;i<num_threads;i++){
38      pthread_join(threads[i],NULL);
39  }
40
41  printf("The sum of the first %d numbers is %d.\n",limit,sum);
42  }
43
44  void *thread_sum(void *args){
45      int i;
46      int *p = (int *)args;
47      int lower = p[0];
48      int upper = p[1];
49
50      for(i=lower+1;i<=upper;i++){
51          pthread_mutex_lock(&m);
52          sum = sum+i;
53          pthread_mutex_unlock(&m);
54      }
55
56      pthread_exit(NULL);
57  }

```

The program below spawns 4 tasks into 3 different threads and then uses *barrier* to wait for them to finish.

```
1 void main(){
2     spawn {
3         /* work */
4         printf("Task 1\n");
5     }
6     spawn {
7         /* work */
8         printf("Task 2\n");
9     }
10    spawn {
11        /* work */
12        printf("Task 3\n");
13    }
14    spawn {
15        /* work */
16        printf("Task 4\n");
17    }
18
19    barrier;
20
21    printf("Done.\n");
22 }
```

The C-equivalent of the SMPL program above:

```
1 #include <stdio.h>
2 #include <stdlib.h>
3 #include <pthread.h>
4
5 void *thread_work(char *);
6
7 void main(){
8     int i;
9     int num_threads;
10
11    printf("Enter number of threads: ");
12    scanf("%d",&num_threads);
13
14    pthread_t threads[num_threads];
15
16    for(i=0;i<num_threads;i++){
17        char str[7];
```

```
18     strncpy(str, "Task ", 5);
19     strcat(str, (i+1), 1);
20     str[6] = '\\0';
21     pthread_create(&threads[i], NULL,
22                  thread_work, (char *)str);
23 }
24
25 for(i=0; i<num_threads; i++){
26     pthread_join(threads[i], NULL);
27 }
28
29 printf("Done!");
30 }
31
32 void *thread_work(char *str){
33     printf("%s", str);
34     pthread_exit(NULL);
35 }
```