

Synchronizing Parallel Tasks Using STM

Ryan Saptarshi Ray, Utpal Kumar Ray, Surojit Mondal, Parama Bhaumik

Junior Research Fellow Department of Information Technology Jadavpur University, Kolkata, India

Associate Professor Department of Information Technology Jadavpur University, Kolkata, India

B.E.(I.T.) Final Year Student Department of Information Technology Jadavpur University, Kolkata, India

Assistant Professor Department of Information Technology Jadavpur University, Kolkata, India

Abstract

The past few years have marked the start of a historic transition from sequential to parallel computation. The necessity to write parallel programs is increasing as systems are getting more complex while processor speed increases are slowing down. Current parallel programming uses low-level programming constructs like threads and explicit synchronization using locks to coordinate thread execution. Parallel programs written with these constructs are difficult to design, program and debug. Also locks have many drawbacks which make them a suboptimal solution. One such drawback is that locks should be only used to enclose the critical section of the parallel-processing code. If locks are used to enclose the entire code then the performance of the code drastically decreases.

Software Transactional Memory (STM) is a promising new approach to programming shared-memory parallel processors. It is a concurrency control mechanism that is widely considered to be easier to use by programmers than locking. It allows portions of a program to execute in isolation, without regard to other, concurrently executing tasks. A programmer can reason about the correctness of code within a transaction and need not worry about complex interactions with other, concurrently executing parts of the program. If STM is used to enclose the entire code then the performance of the code is the same as that of the code in which STM is used to enclose the critical section only and is far better than code in which locks have been used to enclose the entire code. So STM is easier to use than locks as critical section does not need to be identified in case of STM.

This paper shows the concept of writing code using Software Transactional Memory (STM) and the performance comparison of codes using locks with those using STM. It also shows why the use of STM in parallel-processing code is better than the use of locks.

Keywords- Parallel Programming; Multiprocessing; Locks; Transactions; Software Transactional Memory

I. INTRODUCTION

Generally one has the idea that a program will run faster if one buys a next-generation processor. But currently that is not the case. While the next-generation chip will have more CPUs, each individual CPU will be no faster than the previous year's model. If one wants programs to run faster, one must learn to write parallel programs as now multi-core processors are becoming more and more popular. The past few years have marked the start of a historic transition from sequential to parallel computation. The necessity to write parallel programs is increasing as systems are getting more complex while processor speed increases are slowing down. Parallel Programming means using multiple computing resources like processors for programming so that the time required to perform computations is reduced [1].

II. BUS TICKET COUNTER PROBLEM

In the Bus Ticket Counter Problem initially only one counter is open from which passengers may purchase tickets. As more counters open the options

for the passengers (counters from which they can purchase tickets) increases, hence the time taken for purchasing tickets decreases. The problem is to synchronize the operations of the different counters so that the passengers do not have to face any delay.

III. BUS TICKET COUNTER PROBLEM USING LOCKS

The hardest problem that should be overcome when writing parallel programs is that of synchronization. Multiple threads may need to access the same locations in memory and if careful measures are not taken the result can be disastrous. If two threads try to modify the same variable(s) at the same time, data can become corrupt. Currently locks are used to solve this problem. Locks ensure that a critical section, which is a block of code that contains variable(s) that may be accessed by multiple threads, can only be accessed by one thread at a time. When a thread tries to enter a critical section, it must first acquire that section's lock. If another thread is already holding the lock, the former thread must wait until the

lock-holding thread releases the lock, which it does when it leaves the critical section [2].

In the parallel program using threads and locks which solves the Bus ticket-counter problem the time taken for processing of passenger's request at each counter(x) and the number of passengers at each counter(y) are taken as input. There is one thread function-“tctr()”.

The following code snippet shows the tctr thread:

```
void *tctr(int *num_ptr)
{
    unsigned long j;
    int num,*number_ptr;

    number_ptr=num_ptr;
    num=*number_ptr;
    pthread_mutex_lock(&mutex1);
    arr[num]=x*y;

    if(arr[num]<=proc)
    {
        proc= arr[num];
        if(x<prev_bestproc)
        {
            prev_bestproc=x;
        }
    }

    pthread_mutex_unlock(&mutex1);
    pthread_exit(0);
}
```

In the thread “tctr” the amount of time for which the passenger has to wait to purchase ticket from that counter is calculated by the following code snippet:

```
pthread_mutex_lock(&mutex1);
arr[num]=x*y;

if(arr[num]<=proc)
{
    proc= arr[num];
    if(x<prev_bestproc)
    {
        prev_bestproc=x;
    }
}

pthread_mutex_unlock(&mutex1);
pthread_exit(0);
```

3 lock calls are being used in the program. pthread_mutex_init(&mutex1,NULL), is used for lock initialization. pthread_mutex_lock(&mutex1), is used for locking. This means that any thread which needs to access the critical section has to first acquire the lock on mutex1.

pthread_mutex_unlock(&mutex1), is used for unlocking.

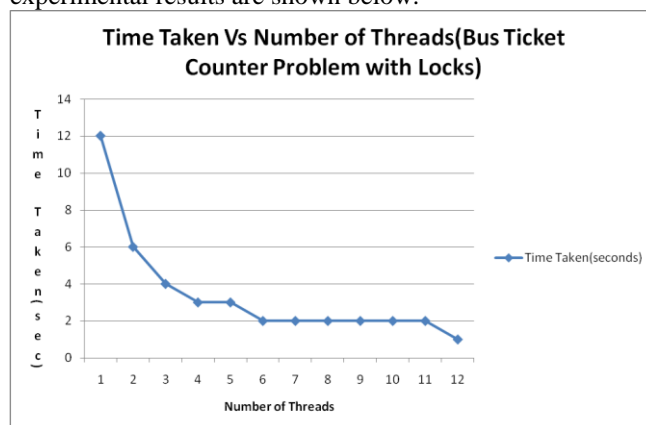
In the program the region where more than one thread may access the global array arr at the same time is the critical section. Thus this region is enclosed within locks. Hence there is no synchronization problem in the above code.

IV. EXPERIMENTAL RESULTS FOR BUS TICKET-COUNTER PROBLEM USING LOCKS

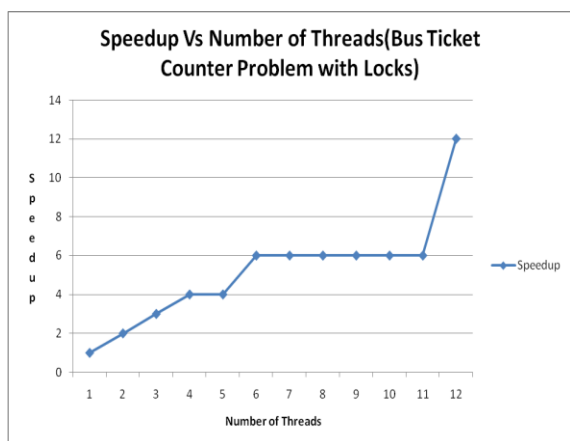
The following table shows the experimental results for Bus ticket-counter problem using locks:

NUMBER OF THREADS	TIME TAKEN(SEC)	SPEEDUP	EFFICIENCY
1	12	1	1.00
2	6	2	1.00
3	4	3	1.00
4	3	4	1.00
5	3	4	0.80
6	2	6	1.00
7	2	6	0.86
8	2	6	0.75
9	2	6	0.67
10	2	6	0.60
11	2	6	0.55
12	1	12	1.00

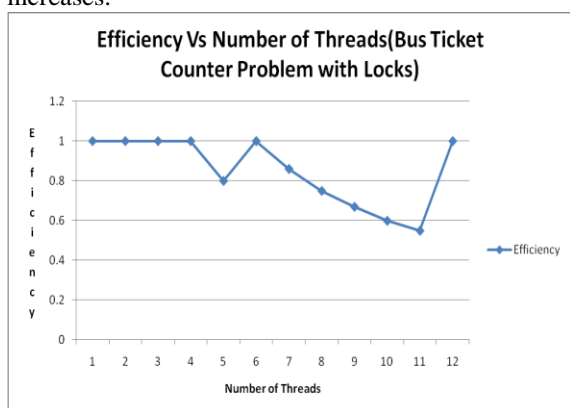
The corresponding graphs for the above experimental results are shown below:



From the above graph we can see that as the number of threads increases the time taken decreases.



From the above graph we can see that as the number of threads increases the speedup also steadily increases.



From the above graph we can see that as the number of threads increases the efficiency varies around 1.

V. BUS TICKET-COUNTER PROBLEM USING STM

The synchronization problem can also be solved using STM. If STM is used in a program then we do not have to use locks in the program. Thus the problems which occur due to the presence of locks in a program do not occur in this type of code. The critical section of the program has to be enclosed within a transaction. Then STM by its internal constructs ensures synchronization in the program.

The structure of the program using threads and STM which solves the bus ticket-counter problem is same as that of the program using threads and locks. The only difference is that STM is being used in this program.

The following code snippet shows the tctr thread:

```
void *tctr(int *num_ptr)
{
    unsigned long j;
    unsigned char byte_under_stm;
    int num,*number_ptr;
```

```
    number_ptr=num_ptr;
    num=*number_ptr;

    stm_init_thread();

    START(0,RW);
    byte_under_stm=(unsigned char)
LOAD(&arr[num]);
    byte_under_stm=x*y;

    if(byte_under_stm<=proc)
    {
        proc= byte_under_stm;
        if(x<prev_bestproc)
            prev_bestproc=x;
    }

    STORE(&arr[num],byte_under_stm);
    COMMIT;
    stm_exit_thread();

    pthread_exit(0);
}
```

The STM functions and calls which have been used in the code are explained below:

stm_init is used to initialize the TinySTM library at the outset. It is called from the main thread before accessing any other functions of the TinySTM library.

stm_init_thread is used to initialize each thread that will perform transactions. It is called once from each thread that performs transactional operations before the thread calls any other functions of the TinySTM library. In this program it is called from the thread tctr.

stm_exit is the corresponding shutdown function for stm_init. It cleans up the TinySTM library. It is called once from the main thread after all transactional threads have completed execution.

stm_exit_thread is the corresponding shutdown function for stm_init_thread. It cleans up the transactional thread. It is called once from each thread that performs transactional operations upon exit. In this program it cleans up the thread tctr.

START(0,RW) is used to start a transaction. In this program it is used in the thread tctr.

COMMIT is used to close the transaction. In this program it is used in the thread tctr.

byte_under_stm=(unsigned char) LOAD(&arr[num]);

stores the value of arr[num] in byte_under_stm. In this program it is used in the thread tctr.

STORE(&arr[num],byte_under_stm);
 stores the value of byte_under_stm in arr[num]. In this program it is used in the thread tctr.

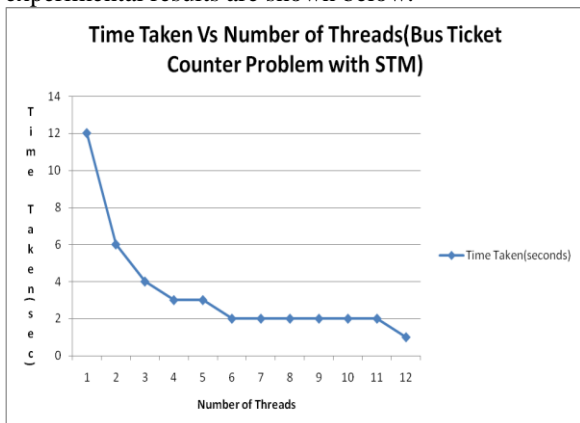
In this program the region where more than one thread may access the global array arr at the same time is the critical section. Thus this region is enclosed within transaction using TinySTM which is a type of STM. Hence there is no synchronization problem in the above code.

VI. EXPERIMENTAL RESULTS FOR BUS TICKET-COUNTER PROBLEM USING STM

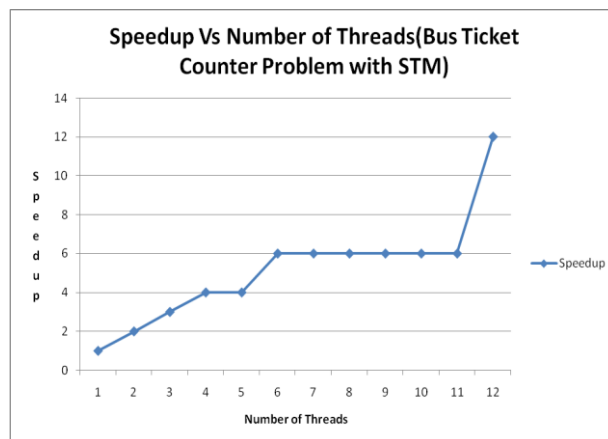
The following table shows the experimental results for bus ticket-counter problem using STM:

NUMBER OF THREADS	TIME TAKEN (SEC)	SPEEDUP	EFFICIENCY
1	12	1	1.00
2	6	2	1.00
3	4	3	1.00
4	3	4	1.00
5	3	4	0.80
6	2	6	1.00
7	2	6	0.86
8	2	6	0.75
9	2	6	0.67
10	2	6	0.60
11	2	6	0.55
12	1	12	1.00

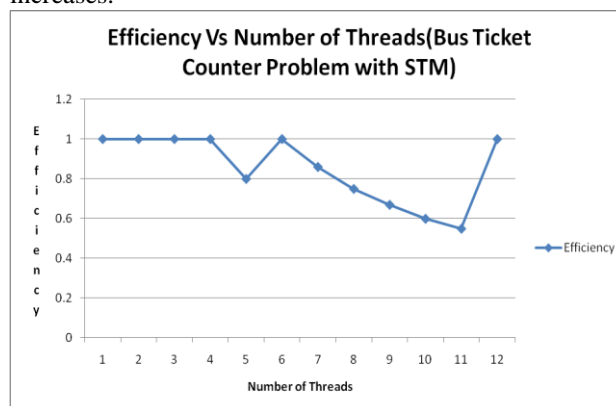
The corresponding graphs for the above experimental results are shown below:



From the above graph we can see that as the number of threads increases the time taken decreases.



From the above graph we can see that as the number of threads increases the speedup also steadily increases.



From the above graph we can see that as the number of threads increases the efficiency varies around 1.

VII. PERFORMANCE COMPARISON OF LOCKS AND STM

From the above experimental results we see that performance of locks and STM are similar.

In the code with locks we have enclosed only the critical section with locks. When we enclosed the entire code with locks then the performance drastically decreased. In the code with STM also we have enclosed only the critical section with STM. When we enclosed the entire code with STM then also the performance remained same. So it can be said that performance of STM is better than that of locks. Also we can say that STM is easier to use than locks as critical section need not be identified in case of STM.

VIII. CONCLUSION

STM has been shown in many ways to be a good alternative to locks for writing parallel programs. STM provides a time-tested model for isolating concurrent computations from each other. This model

raises the level of abstraction for reasoning about concurrent tasks and helps avoid many parallel programming errors.

This paper has discussed how STM can be used to solve the problem of synchronization in parallel programs. STM has ensured that lock-free parallel programs can be written. This ensures that the problems which occur due to the presence of locks in a program do not occur in this type of code. It has also been shown that STM is easier to use than locks as critical section need not be identified explicitly in case of STM. In case of STM if the entire code is enclosed within STM the performance of the code is same as that of the code in which only the critical section is enclosed within STM. But in case of locks if the entire code is enclosed within locks then the performance decreases sharply. So it has been shown that the performance of STM is much better than that of locks.

Many aspects of the semantics and implementation of STM are still the subject of active research. While it may still take some time to overcome the various drawbacks, the necessity for better parallel programming solutions will drive the eventual adoption of STM. Once the adoption of STM begins it will pick up momentum and make a very large impact on software development in the long run. In the near future STM will become a central pillar of parallel programming.

REFERENCES

- [1] Simon Peyton Jones, "Beautiful concurrency".
- [2] Elan Dubrofsky, "A Survey Paper on Transactional Memory".
- [3] Pascal Felber, Christof Fetzer, Torvald Riegel, "Dynamic Performance Tuning of Word-Based Software Transactional Memory".
- [4] http://en.wikipedia.org/wiki/Transactional_memory
- [5] James Larus and Christos Kozyrakis. "Transactional Memory"
- [6] Pascal Felber, Christof Fetzer, Patrick Marlier, Torvald Riegel, "Time-Based Software Transactional Memory"
- [7] Tim Harris, James Larus, Ravi Rajwar, "Transactional Memory"
- [8] Mathias Payer, Thomas R. Gross, "Performance Evaluation of Adaptivity in Software Transactional Memory"
- [9] Kevin E. Moore, Jayaram Bobba, Michelle J. Moravan, Mark D. Hill, David A. Wood., "LogTM: Log-based Transactional Memory"
- [10] Dave Dice , Ori Shalev , Nir Shavit., "Transactional Locking II"
- [11] <http://tmware.org>
- [12] Maurice Herlihy, J. Eliot B. Moss, "Transactional Memory: Architectural Support for Lock-Free Data Structures".
- [13] Martin Schindewolf, Albert Cohen, Wolfgang Karl, Andrea Marongiu, Luca Benini, "Towards Transactional Memory Support for GCC".
- [14] Virendra J. Marathe, Michael F. Spear, Christopher Heriot, Athul Acharya, David Eisenstat, William N. Scherer III, Michael L. Scott, "Lowering the Overhead of Nonblocking Software Transactional Memory".
- [15] Utku Aydonat, Tarek S. Abdelrahman, Edward S. Rogers Sr., "Serializability of Transactions in Software Transactional Memory".
- [16] Maurice Herlihy, Nir Shavit, "The Art of Multiprocessor Programming".
- [17] Brendan Linn, Chanseok Oh, "G22.2631 project report: software transactional memory".
- [18] Ryan Saptarshi Ray, "Writing Lock-Free Code using Software Transactional Memory".
- [19] Ryan Saptarshi Ray,Utpal Kumar Ray "Different Approaches for improving performance of Software Transactional Memory".