

APPLICABILITY CRITERIA OF THE STATISTICAL SYNCHRONIZATION METHOD

Gábor Lencse

Department of Telecommunications
Technical University of Budapest
Sztoczek utca 2
H-1111 Budapest, Hungary
E-mail: Gabor.Lencse@hit.bme.hu
<http://www.hit.bme.hu/phd/lencse>

Department of Telecommunications
Széchenyi István College
Hédervári u. 3.
H-9026 Győr, Hungary
E-mail: lencse@szif.hu
<http://rs1.szif.hu/~lencse>

KEYWORDS

parallel discrete-event simulation, statistical synchronization, applicability criteria, efficiency, accuracy.

ABSTRACT

The applicability of the Statistical Synchronization Method (SSM) in parallel discrete-event simulation is studied. Criteria are given for the time-driven approach (SSM-T). It is proven, that the level of the output error can be guaranteed. Positive and negative examples are shown to give a deeper insight to the criteria. We conclude that SSM-T can be a very efficient synchronization method in the parallel simulation of large communications and other systems.

INTRODUCTION

The Statistical Synchronization Method (SSM) (Pongor 1992) is a promising alternative to the conventional synchronization methods for parallel discrete event simulation (e.g., conservative, optimistic) (Fujimoto 1990).

The conventional synchronization methods use event-by-event synchronization between the segments of the simulated system and they are unfortunately not applicable to all cases or do not provide the desirable speed-up. SSM does not exchange individual messages between the segments but rather the statistical characteristics of the message flow. Actual messages are regenerated from the statistics at the receiving side. (Further explanation will be given later.) SSM claims to be less sensitive to communication delay and it requires less network bandwidth than event-by-event methods. Nevertheless, it is not accurate in the sense that an event that occurred in one segment of the system does not have an immediate influence on another segment. For this reason, the method cannot be applied in some simulations, for example in the case of digital circuits but remains feasible in other classes of simulation such as the performance estimation of communication systems.

An addition advantage of SSM is that it is relatively easy to extend existing non-parallel simulation software for use with SSM, which is not necessarily true for other synchronization methods. An example of a simulation system that directly supports SSM is OMNeT++¹ (Varga, 1998b).

The aim of this paper is to determine under what conditions the Statistical Synchronization Method can be applied in parallel discrete-event simulation. We are concerned with both the accuracy of the results and the speed-up can be achieved.

The remainder of this paper is organized as follows: after a brief introduction to SSM and SSM-T, the applicability criteria for SSM-T are given in an informal way. They are formalized in the next chapter together with a proof that the required level of the output error can be guaranteed by satisfying the criteria. Afterwards, we show positive and negative examples to give a better insight of the criteria. Finally, we present the conditions for a good speed-up and give the directions of future research on the statistics exchange control algorithm of SSM-T.

This topic was identified as being of importance in the parallel simulation of large (communication) systems.

THE STATISTICAL SYNCHRONISATION METHOD

The Original SSM

For those not familiar with SSM, a short summary of the Statistical Synchronization Method is given here. See (Pongor 1992) for more information.

Similarly to other parallel discrete-event simulation methods, the model to be simulated — which is more or less a precise representation of a real system — is divided into segments, where the segments usually describe the behavior of

¹ OMNeT++ is freely available from the Web site cited in the Reference

functional units of the real system. The communication of the segments can be represented by sending and receiving various messages. For SSM, each segment is equipped with one or more input and output interfaces. The messages generated in a given segment and to be processed in a different segment are not transmitted there but the *output interfaces* (OIF) collect statistical data of them. The *input interfaces* (IIF) generate messages for the segments according to the statistical characteristics of the messages collected by the proper output interfaces. (See Figure 1.)




Figure 1. An OIF - IIF pair

The segments with their input and output interfaces can be simulated separately on separate processors, giving statistically correct results. The events in one segment have not the same effect in other segments as in the original model, so the results collected during SSM are not exact. The precision depends on the partitioning of the model, on the accuracy of statistics collection and regeneration, and on the frequency of the statistics exchange among the processors.

SSM-T: Refinement Of SSM

The original SSM does not explicitly state when the OIF's should send their statistics to the appropriate IIF's. The results of (Lencse 1998b) would suggest that the statistics collection must be continued until the sample contains the required number of observations, then the statistics should be sent and the statistics collection should be restarted. This was called *SSM-C* (the counter driven approach) in (Lencse 1998a). In that paper, *SSM-T* (the time driven approach) was proposed for parallel simulation, which works as follows. Using the notations of Figure 1, at the beginning of the simulation the OIF of segment A must tell the IIF of segment B at what virtual time it will send the first statistics. This is t_1 in Figure 2. In this figure the thin horizontal lines show the wall-clock (real) time of the processors executing the segments and the thick lines are the virtual times of the segments. Segment B takes into consideration the first statistics from segment A at t_1 virtual time. It is done in the following way: In the figure, segment B receives the first statistics from segment A at t_x (according to its own virtual time) and as $t_x < t_1$, segment B schedules the arrival of the statistics for t_1 . The other possibility is shown on the example of t_2 . Segment B has not received the statistics until t_2 , and it has no more events scheduled with less than or equal time stamp, so the execution of the simulation of segment B is suspended until the statistics arrive from segment A. Then segment B receives the statistics and the execution resumes. Segment B always knows at what virtual time to expect the next statistics, because Segment A always includes the virtual arrival (=sending) time of

the i -th statistics in the $(i-1)$ th statistics package. We called this solution *loose synchronisation*. This method makes it possible for the simulation of the segments to run independently on separate processors in the vast majority of time and therefore it may result in excellent speed-up. See (Lencse 1998a) for more details.




Figure 2. The operation of SSM-T. See the text for explanation.

THE APPLICABILITY CRITERIA OF SSM-T

SSM-T can be applied and produces meaningful results if the following conditions are met: (The expressions written in italic will be formally defined in the next chapter.)

- The simulated system can be divided into segments so that not the individual messages but only their statistical characteristics are important between the segments.
- A *small error* in the approximation of the statistical characteristics of the message flow causes *small error* in the output of the simulation that *depends only on the measure of the approximation error*.
- The parameters of the model may change during the simulation but the changes in the statistical characteristics of the message flows between the segments are *rare enough*.
- Note that a change in the statistical characteristics of the message flow is only propagated to other segments when the statistics package is actually sent over by the OIF. The fourth condition is that this delay causes error in the output of the simulation only during the delay and/or at most during an additional time interval with a length proportional to the delay.

THE OUTPUT ERROR OF SIMULATION WITH SSM-T

Let us denote the concerned statistical characteristics of the message flow by the random variable X and its approximation by X^* . The error of the approximation is also a random variable: $h_X = X^* - X$. The observed output of the simulation is denoted by O . The h_X error of the approximation causes $h_O = O^* - O$ error in the output. Condition (b) is defined formally as follows:

- $h_O = f(h_X)$ for some f , and $\forall \varepsilon > 0 \exists \delta: |h_X| < \delta \Rightarrow |h_O| < \varepsilon$.

Note, that sometimes the following conditions may be enough:

(b'') $h_o = f(h_X)$ for some f , and $\forall \epsilon > 0 \exists \delta: E\{h_X\} < \delta \Rightarrow E\{h_o\} < \epsilon$, where $E\{\cdot\}$ is the expected value of the random variable.

Let us denote the number of observations in a sample by n . $\forall \delta > 0 \exists N: n > N \Rightarrow |h_X| < \delta$. The value of N depends on both the required value of δ and the convergence speed of the statistics collection method used. See (Lencse 1998b) for more details about the convergence speed of some well-know statistics collection methods.

In the stationary case we can guarantee $|h_o| < \epsilon$ with the appropriate choice of N .

Let us consider the transients in the system. If the distribution of the random variable X changes in segment A at t_c in the i -th sample collection interval, the exact new statistics arrive at segment B at the t_{i+1} synchronisation point of virtual time. (see Figure 3) If T_N is an upper bound for the length of any $(t_{i+1} - t_i)$ interval, then the length of the transient $T_{TR} < 2T_N$.

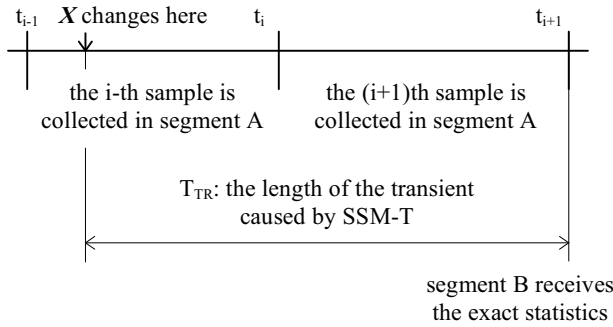


Figure 3. The transient caused by SSM-T

Condition (d) says that the transient may cause output error during the T_{TR} time of the transient, plus maximum $c_e T_{TR}$ time after it. (c_e is an appropriate constant.) Thus the output of the simulation may contain an error due to the transient caused by SSM-T less than $c_e T_N$ time, where $c_e = 2(1+c_e)$.

Let us denote the time elapses from the end of the $c_e T_N$ time interval to the next change of X by T_{QST} , the time of the quasi stationary state. The empirical meaning of "rare enough" in condition (b) is: $c_e T_N \ll T_{QST}$. Let us make it a bit more formal: H_o is an upper bound for the absolute value of the output error during the transient caused by SSM-T. Now, we show that $|\bar{h}_o| < \epsilon$ can be ensured. Let us choose δ' : $|h_X| < \delta' \Rightarrow |h_o| < \frac{1}{2}\epsilon$ during the T_{QST} time period. Let N_{TR} and N_{QST} denote the number of the collected output statistics values during the $c_e T_N$ and the T_{QST} time periods, respectively. The before mentioned "rare enough" condition is:

$$(c') \quad N_{QST} \geq \frac{2N_{TR}H_o}{\epsilon} - N_{TR}$$

The average output error is:

$$|\bar{h}_o| \leq \left| \frac{N_{TR}H_o}{N_{QST} + N_{TR}} \right| + \left| \frac{N_{QST}h_o}{N_{QST} + N_{TR}} \right| \leq \frac{1}{2}\epsilon + \frac{1}{2}\epsilon \leq \epsilon$$

However, the condition for N_{QST} is quite strict, and it is not always necessary. We have eliminated the error of the output of the simulation by averaging very many values. If we know the time of the changes of X in advance, or if we can detect the change, we can just omit the N_{TR} number of output statistics values with error. By doing so, a lot of virtual time (and therefore simulation execution time too) may be saved, because in this way only the $c_e T_N$ virtual time is wasted, and the requirement for N_{QST} is just the same as it is in the case of the traditional event-by event synchronisation:

(c'') N_{QST} must be large enough to collect the statistics of the output of the simulation with the required accuracy.

The proportion of the wasted $c_e T_N$ and useful T_{QST} virtual time is very important. The $c_e T_N$ virtual time is used up just to eliminate the transient caused by SSM-T. In addition to its execution time comes the execution time wasted due to communication overhead between the processors executing the segments. However, if both of them are low compared to the execution time of T_{QST} and the model of the simulated system was partitioned in the way that the load of the executing processors is nearly balanced, our simulation may produce a good speed-up.

EXAMPLES FOR THE APPLICATION OF SSM-T

Satellite Power Consumption - A Positive Example

Let us consider the following example. A hurricane forecasting satellite collects data of the atmosphere and after some preprocessing, it sends them to the Earth for evaluation. The whole system is built up by three major functional units:

1. The Intelligent Measurement System controls the sensors and evaluates their signals, collects and preprocesses measurement data. Its output is a variable rate packet data flow. The packet rate depends on environmental conditions such as the state of the atmosphere, hurricane suspicious observations, etc.

2. The Data Transmission System carries the data packets from the Intelligent Measurement System to the Data Evaluation System on the Earth. The Data Transmission System contains a radio link downwards and another one upwards. The data sent through the downlink is acknowledged on the uplink. The transmission power is controlled in the function of the packet loss ratio, so the power consumption/bit depends on the environmental conditions (state of the atmosphere, orbit deviations, etc.) too. To save power, the carrier of the downlink is turned off when there is no transmission.

3. The Earth Data Evaluation System is responsible for the final evaluation of the collected measurement data.

Figure 4 shows the block diagram of the physical system.

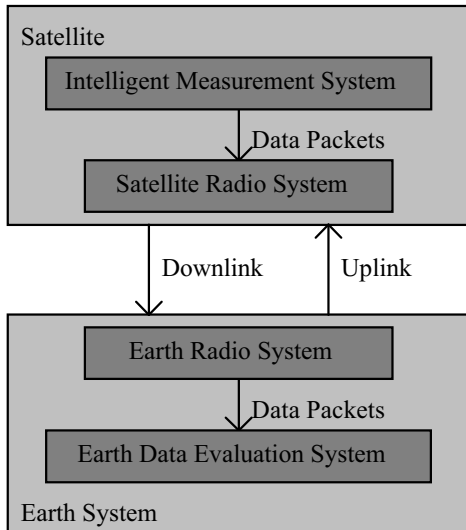


Figure 4. The satellite data collection system

Our task is to determine the behavior of the power consumption of the Satellite Radio System to be able to tell the solar cell and battery requirements. The power consumption depends on both the packet rate from the Intelligent Measurement System and the radio channel conditions. Both quantities depend on environmental conditions that are too complex for an analytical treatment. Some of the environmental conditions (e.g. atmosphere) influence both the packet rate and the required transmission power, so they cannot be simulated independently. However, it is known from earlier experiments that the environmental conditions change very slowly compared to the data packet rate, that is typically millions of data packets are transmitted between two consecutive significant changes in the environmental conditions. It is also known that the channel capacity is more than twice as much than it is necessary for the maximum packet rate, so there is practically not buffering except that the packets are stored (to be able to retransmit them) until an acknowledgement is received. We propose the following parallel simulation:

- The Earth Data Evaluation System is omitted as it has no influence for the investigated power consumption.
- The remainder of the system is divided into two segments: the Intelligent Measurement System and the Data Transmission System, they are simulated parallel on two processors. The segments model those environmental conditions that are relevant to their operation.
- The two segments of the model are executed by two processors. SSM-T is applied between the two segments. The packet inter-arrival time statistics are collected by the OIF of the Intelligent Measurement System and the result is sent to the IIF of the Data Transmission System. The packet data flow is regenerated by the IIF of the Data Transmission System.

Figure 5 shows the block diagram of the simulation model.




Figure 5. The simulation model of the satellite data collection system

The applicability criteria of the SSM-T are satisfied:

- At the boundaries of the two segments, not the individual packets but only the average packet rate is important in the point of view of the power consumption.
- A small error of the estimation of the packet rate (or the distribution of the packet inter-arrival time) causes a small error in the calculation of the power consumption and depends only on the measure of the error, not the actual value of the packet rate.
- Significant changes are rare enough to make the necessary number of observations in quasi stationary state.
- As there are practically no buffering, the delay of the changes of the packet rate causes error only until the arrival of the new perfect statistics.

Negative Examples

The first example is the simulation of an FDDI (ANSI X3.139, 1987) ring. If the ring is divided into two (or more) segments and the segments are simulated by separate processors, the explicit passing of the token cannot be replaced by the arrival time statistics of the token collected at the segment boundaries. This would be a violation of the Media Access Control protocol resulting in ring recovery (token loss, duplicate tokens etc.). Here SSM-T cannot be applied, because condition (a) is not satisfied.

Our second negative example is derived from the before mentioned positive one. Let us modify the system described there in the following way: The Data Transmission System contains two downlinks with no transmission power control. The unacknowledged packets are retransmitted until an acknowledgment arrives for the packet. The primary downlink is always operational and the secondary works only if it is necessary due to high packet rate and poor channel conditions (that is, a high number of packet retransmissions). The earth station does not have burst demodulators, so when the second carrier is put into operation, there is a significant synchronization overhead. For this reason, the carrier is not switched off immediately after transmitting a packet, but only after a certain delay. This also means that even a few packets on

the secondary link may result in significantly increased power consumption.

For any given channel conditions one can calculate what packet rate can completely exhaust the capacity of the primary downlink. An arbitrarily small error in the estimation of the packet rate from the Intelligent Measurement System can cause a serious error in the output of the simulation if the packet rate is close enough to the calculated critical rate. Now, condition (b) is not satisfied.

Two Interconnected FDDI Rings - Another Positive Example

In (Lencse 1998a) SSM-T was used in the simulation of the FDDI backbone of the Technical University of Budapest. This network consists of two rings: The Northern Ring is a university-wide network and consists of 15 FDDI stations interconnected by 5 wiring concentrators. The Southern Ring is the backbone of the Faculty of Electrical Engineering and Informatics, and being smaller ring of 7 FDDI stations. The two rings are interconnected by a router. The topology of the network and the cable lengths were taken from the real system. The load used in the simulation model came from measurements taken on the real FDDI rings.

First, a very accurate simulation of the network was performed. Then, SSM-T interfaces were inserted between the two FDDI rings and the simulation was performed with the same parameters as before. The utilization of the rings were measured in both cases. It was found that the utilization values were close to each other in the two simulations.

Space permits no more detailed discussion of the experiments here, but the reader is encouraged to read the original paper and check if the applicability criteria of SSM-T are really satisfied.

CONDITIONS FOR A GOOD SPEED-UP

In the practical application of SSM-T, the length of the virtual time interval while the OIF collects a statistics package is very important. There are three points:

1. The time interval should be large enough to collect a sample that is based on enough observations to produce an estimation with the required accuracy.
2. The time interval should be short enough to bound the length of the transient caused by SSM-T.
3. The frequency of the statistics exchange should not be too high to produce good speed up.

The first two conditions should be evident for the reader, but the third one requires some explanation. Until now, the simulation with SSM-T was independent from the execution environment. However, the main aim of SSM-T is to produce both good results and good speed-up. Thus, the overhead caused by the statistics exchange is very important. Let us consider this overhead:

The statistics transmission directed graph is defined as follows:

- the nodes are the segments of the simulation model,
- the (directed) edges are the segment to segment routes of statistics transmission.

If there are no directed loops in the graph then the simulation may work as a pipeline with infinite buffers between the stages. However, if there is a directed loop in the graph the virtual times of the segments of this loop are synchronized in some way. Let us consider the simplest example: there are two segments and they send statistics to each other. Let them exchange their statistics every T virtual time interval. The segments are executed by two processors A and B. The processors are identical and they do not have any other load. The execution time of the T virtual time is τ_A and τ_B . The time of communication is denoted by τ_C . This time includes the time of data packing and conversion to the network data format (e.g. XDR, if necessary) data transmission time and propagation delay plus data conversion from network format (if necessary) and unpacking. The execution time of a T virtual time interval with SSM-T is:

$$(1) \tau_2 = \max(\tau_A, \tau_B) + \tau_C$$

The overhead of the statistics collection and regeneration done by the IIF's and OIF's is denoted by τ_{IIF} τ_{OIF} . These are included in τ_A and τ_B , so they must be subtracted in the calculation of the execution time of the traditional simulation. The execution time of a T long virtual time interval using traditional uni-processor simulation is:

$$(2) \tau_1 = \tau_A - \tau_{OIF-A} - \tau_{IIF-A} + \tau_B - \tau_{OIF-B} - \tau_{IIF-B} + \tau_C$$

Let us group the I/O interfaces overhead into τ_{IF} .

$$(3) \tau_{IF} = \tau_{OIF-A} + \tau_{IIF-A} + \tau_{OIF-B} + \tau_{IIF-B}$$

The speed-up is:

$$(3) s = \frac{\tau_A + \tau_B - \tau_{IF}}{\max(\tau_A, \tau_B) + \tau_C}$$

This value can be close to 2 if $\tau_A \approx \tau_B$, $\tau_C \ll \tau_A$ and $\tau_{IF} \ll \tau_A$, that is the load of the processors is well balanced, the communication overhead, and the overhead caused by the statistics collection and message regeneration are small.

In a large (communication) system there are usually multiple points where the SSM-T applicability criteria are satisfied. It means that the insertion of the SSM-T interfaces to these points will not cause significant degradation of the accuracy of the results. Out of these point the simulationist must carefully select those, that separate the model to segments of similar complexity, so that the computing power required by the simulation of the segments are in the same order.

When selecting the statistics collection method it is worth considering its algorithmical complexity (Lencse 1998b). An interesting new density estimation method, k-split (Varga 1998a) may also be used in SSM OIF's.

And last but not least the frequency of the statistics exchange should not be higher than it is required. In earlier experiments on simulating two interconnected FDDI rings by two processors using SSM-T (Lencse 1998a), we achieved 1.75, 1.86 and 1.91 speed-up depending on the frequency of statistics exchange.

The direction of the future research is to develop a statistics exchange control algorithm that determines the optimal time of the statistics exchange.

The existence of weaker criteria for the applicability of SSM-T is also a question of interest.

CONCLUSION

The applicability of the modified Statistical Synchronization Method (SSM-T) was studied. Criteria were given for the applicability of SSM-T in parallel discrete-event simulation. We have proven that the small level of output error of the parallel simulation using SSM-T compared to the uni-processor simulation without SSM-T can be guaranteed if the criteria are satisfied.

We showed a real life example where the applicability criteria are satisfied and SSM-T can be applied. We gave negative examples too.

We presented the conditions when the application of SSM-T results in a good speed-up.

We conclude that SSM-T can be a very efficient synchronization method in the parallel simulation of large (communication) systems.

The direction of the future research is to develop a statistics exchange control algorithm that determines the optimal time of the statistics exchange.

REFERENCES

ANSI X3.139. 1987. Fiber Distributed Data Interface (FDDI) Token Ring Media Access Control (MAC). American National Standards Institute, New York, NY

Fujimoto, R. M. 1990. "Parallel Discrete Event Simulation". *Communications of the ACM* 33, no 10, 31-53

Jefferson, D; B. Beckman; F. Wieland; L. Blume; M. DiLoreto; P. Hontalas; P. Laroche; K. Sturdevant; J. Tupman; V. Warren; J. Vedel; H. Younger and S. Bellenot. 1987. "Distributed Simulation and the Time Warp Operating System". *Proceedings of the 12th SIGOPS - Symposium on Operating System Principles*, 73-93.

Lencse, G. 1997. "Efficient Simulation of Large Systems - Transient Behaviour and Accuracy" *Proceedings of the 1997 European Simulation Symposium (ESS'97)* (Passau, Germany, Oct. 19-23). SCS Europe, 660-665.

Lencse, G. 1998a. "Efficient Parallel Simulation with the Statistical Synchronization Method" *Proceedings of the Communication Networks and Distributed Systems Modeling and Simulation (CNDS'98)* (San Diego, CA. Jan. 11-14). SCS International, 3-8.

Lencse, G. 1998b. "Statistics Collection for the Statistical Synchronisation Method" *Proceedings of the 1998 European Simulation Symposium (ESS'98)* (Nottingham, UK. Oct. 26-28). SCS Europe, 46-51.

Pongor, Gy. 1992. "Statistical Synchronization: a Different Approach of Parallel Discrete Event Simulation". *Proceedings of the 1992 European Simulation Symposium (ESS 92)* (Nov. 5-8, 1992, The Blockhaus, Dresden, Germany.) SCS Europe, 125-129.

Varga, A. 1998a. "K-split – On-Line Density Estimation for Simulation Result Collection". *Proceedings of the European Simulation Symposium* (Nottingham, UK. Oct. 26-28). SCS Europe, 41-45.

Varga, A. 1998b. "OMNeT++ Discrete Event Simulation System" <http://www.hit.bme.hu/phd/vargaa/omnetpp.htm>

BIOGRAPHY

Gábor Lencse was born in Győr, Hungary, in 1970. He received his M.S. in electrical engineering and computer systems from the Technical University of Budapest in 1994. He is currently pursuing his Ph. D. at the same university. The area of his research is computer architectures and parallel processing. He is interested in (parallel) discrete event simulation. Since 1997, he works for the Széchenyi István College in Győr. He teaches computer networks and networking applications running over TCP/IP. He is a member of the Society for Computer Simulation International.