
Discrete Event Simulation in Forth

Leonard F. Zettel

Room S3035, SRL

Ford Motor Company

P.O. Box 2053

Dearborn, Michigan 48121, U. S. A.

Abstract

A Forth based discrete event simulation system is presented, with words for scheduling events, executing the simulation, or optionally stepping through it. An example problem, the barbershop simulation, is presented and discussed. Forth is shown to have considerable advantage over implementation in conventional procedural languages.

Introduction

Discrete event simulation is a widespread application of computing. This paper will describe the implementation of a discrete event simulation in Forth and demonstrate that, the characteristics of Forth lead naturally to a system with considerable advantages over systems coded in conventional procedural languages like FORTRAN or Pascal. For a thorough description of a FORTRAN based simulation system see [PRI74]. Discrete event simulation has been used to model and study a great variety of phenomena. It is particularly well suited to the study of problems involving waiting lines—queueing problems. Simulation studies have been made of waiting lines at banks [PRI74], the flow of work through factories [SHE84], and the performance of computer systems [SIG84], to mention just a few examples.

A Simulation Problem

To guide the following discussion, let us consider a simple queueing situation, the barber shop problem [SCH74]. Customers arrive at random times, wanting to get a haircut. There is one barber, and (s)he can serve a typical customer in about sixteen minutes. A customer who arrives when the barber is busy takes a seat in the shop and waits his turn. In more technical language, the customer has joined a single-server, first-in, first-out (FIFO) *queue*. A customer arriving when three or more people are already waiting may *balk*—decide not to join the queue but instead get a haircut at another time or place. Items that might be of interest in studying this system, for different arrival rates of customers and haircut times, could include the proportion of time the barber is busy (his *utilization*), the length of time customers have waited (average, maximum, percentile distribution), the length of the waiting line, and the proportion of balks.

In approaching the discrete event simulation problem we consider that at any time the system can be in only one of a set of *states*. For our purposes the state of the barber shop is completely defined by whether the barber is busy or idle and, if busy, by the number of customers waiting to get haircuts. Transitions from state to state occur instantaneously and are called *events*. For a more rigorous definition of a discrete event simulation system see [SUR83]. For the barbershop system we need two events: arrival of a customer, and completion of a haircut.

The commonest approach to discrete event simulation, and the one we will take here, involves the maintenance of a calendar of events. The simulation proceeds by having the system remove the next most imminent event from the calendar, update a simulation clock to the time of the event removed, change the system state according to the rules for that event (which may be conditional on the state of the system at the time of event occurrence), and possibly add one or more new events to the calendar.

In the barber shop problem, the simulation would be started off by showing no one waiting for a haircut, setting the status of the barber to "idle", and placing a customer arrival event on the calendar. When executed, a customer arrival event would do the following:

- 1) Schedule the next customer arrival event.

- 2) If the barber is idle, set the state of the barber to busy and put a haircut completion event on the calendar.

- 3) If the barber is busy and there are fewer than n customers waiting, increase the number of customers waiting by one.

- 4) If there are n or more customers waiting, increase the number of balks by one and do nothing else, thus effectively the customer does not enter the barbershop.

The haircut completion event actions are:

- 1) If there are no customers waiting, set the state of the barber to idle.

- 2) If there are customers waiting, reduce the number of waiting customers by one and schedule another haircut completion event.

To round out a practical system, there would probably be one or more statistics reporting events and an end-of-simulation event, either scheduled to occur at a definite time (eight hours after the shop opened, say) or triggered by some system statistic like the completion of 100 haircuts.

An event calendar is an instance of a priority queue, the priority in this case being the event occurrence time. Schemes for priority queues and event calendars have included linear lists, indexed lists, heaps and leftist trees. Any of these can be implemented in Forth. For a review of these structures and issues related to their relative performance see [MCC81].

In FORTRAN the event type is usually identified by an integer carried as part of the event notice on the calendar. A particular event is executed by means of a subroutine call after a computed GO TO, with all the overhead that that implies. Space for the event calendar and other files, and variables and arrays for system state representation and statistics collection, are kept in large COMMON blocks. Exceeding any of the storage limits, besides being difficult to detect, necessitates often cumbersome and complicated reallocation of memory. As a result, initial allocations err on the generous side, resulting in the waste of computer resources as we wait for the rare worst case to happen.

Simulation in Forth

Proper implementation in Forth avoids all of this. Variables and arrays representing the system state are naturally global. Data structures that will change size during the simulation can dynamically allocate and free space in a manner that does not interfere with other structures, so that the only limit encountered will be that of the overall space accessible to the Forth dictionary. The key data structure is the event calendar, here a linear linked list [KNU72] with variable node size. We use a linear linked list for simplicity, and because it will give performance equal to or better than the alternatives as long as the event calendar does not grow too long. Figure 1 shows a conceptualization of the list, and Figure 2 represents more nearly the way it would appear in the computer memory allocated to the Forth dictionary.

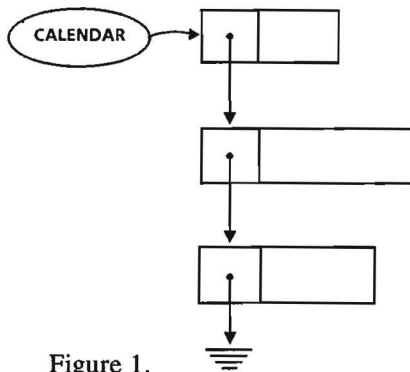


Figure 1.




Figure 2.

The variable **CALENDAR** points to the top of the list. The minimum node size for an event notice is five cells. The first cell contains the pointer to the next node on the list, or is zero if this is the last node. Putting the link node first increases the speed with which we can pass from node to node during processing and also simplifies the code of the list processing routines **PUSH** and **POP**, since they then don't have to distinguish between an occurrence of the address of the pointer variable and the occurrence of a node address. The second cell of a node contains the node size in cells, including the pointer cell and node size cell.

The third and fourth cells contain the double number event time. The fifth cell has the compilation address of the routine associated with the event. If the event routine requires parameters, then they are placed in the sixth and succeeding cells of the node. When the event routine executes the parameters will be in order on top of the parameter stack. The event calendar is arranged in order of increasing execution times, so the notice for the next event to be executed is always at the top of the list.

The event time is expressed as an integer, the customary form of number representation in Forth. In recent years it has become fashionable among some in the simulation community to consider a floating point time an advantage, because its reduced granularity would hopefully eliminate problems which can arise when two distinct events are scheduled for the same moment in time. If the event that ends up first on the calendar causes state changes that influence the action of the second event, then results of the simulation will depend on whether it appears first or second. In the opinion of the author, this "advantage" is likely to be more than offset by the changing granularity of a floating point clock as the simulation progresses. At large simulation times the results of events close to each other in time will be modeled with less and less accuracy. Finally, we note that by using a double number simulation time, a time unit of one second (for many simulations a rather fine granularity) lets a simulation span sixty years before clock overflow.

The Forth word **SCHEDULE** (screen 39) inserts a new event notice at the proper spot in the event calendar. **NEXT.EVENT** (screen 40) takes the top event notice off the calendar, updates the simulation clock (kept in the double number variable **CLOCK**), puts any event attributes on the parameter stack, and **EXECUTE**'s the associated event routine. Here we take advantage of the threaded code properties of Forth to minimize the overhead associated with invoking an event routine. If there is no event notice on the calendar, **NEXT.EVENT** aborts with a message to that effect. The continuing execution of the simulation is accomplished by **SIMULATE**, (screen 40). It simply loops endlessly, invoking **NEXT.EVENT** until interrupted by input from the terminal.

At this point note that implementation in Forth has given us an important bonus at next to no cost. By invoking **NEXT.EVENT** directly from the terminal we can step through the simulation event by event, examining or changing any aspect of the system at our leisure, directly from the keyboard. This is a greater degree of interactive capability than that offered by any other simulation system known to the author.

It is fairly common to want an event routine to schedule a future instance of itself, so we include the word **ME** which will put the compilation address of the word in which it is embedded on the stack

when executed. The definition of **ME** uses words not found in the Forth-79 standard. The usage here is FIG Forth and follows [PET81]. We note in passing that the commonly used **MYSELF** (**RECURSE** in some implementations) can be defined as follows:

```
: MYSELF [COMPILE] ME EXECUTE ; IMMEDIATE
```

Although the author has not yet encountered the situation in practice, it is conceivable that there might be a closed circle of event invocations in a simulation, i.e. event A schedules an instance of routine B which in turn schedules A. This and its more complicated variations can be handled by any of the vectored execution techniques available to the Forth programmer.

Now let us consider some of the specifics of simulating our example (screens 88-90). We define the Boolean constants **TRUE** and **FALSE** according to the Forth-79 standard (in Forth-83 **TRUE** should be **-1**). Variable **WAITING** holds the number of customers waiting to get a haircut. **BALKS** holds the number of customers who have balked since the simulation started. **BARBER** holds the value **TRUE** when the barber is busy, and **FALSE** if (s)he is idle. The arrival time of the next customer will be determined by a random draw of an integer between **ARRIVAL.MIN** and **ARRIVAL.MIN** plus **ARRIVAL.SPREAD** minus one, each integer having an equal probability of being chosen. **HAIRCUT.MIN** and **HAIRCUT.SPREAD** function the same way for haircut completion times.

For convenience we have included the word **SETUP.BARBERSHOP** to initialize the simulation. It sets the number served, the number who have balked, and the simulation time all to zero. The state of the system is put to the traditional "empty and idle" by putting zero in **WAITING** and **FALSE** in **BARBER**. Finally, **CUSTOMER.ARRIVES** is executed.

CUSTOMER.ARRIVES schedules the occurrence of the next customer arrival. The code for **CHOOSE** and **RANDOM** is taken from [BRO81] (screen 17).

Next, if the barber is busy and less than four people are waiting we add one to **WAITING** and exit. Note that since for haircutting purposes all customers are created equal, in this case there is no need to create an actual file of waiting customers. The number waiting is sufficient for all the purposes of the simulation. If the barber is busy and four or more people are waiting, we add one to **BALKS** and exit. If the barber is idle we change the barber state to busy and schedule **FINISH.HAIRCUT**. Note that in this version of Forth, ['] puts the parameter field address of the following word on the stack, so we get the compilation address required by **SCHEDULE** by following the event routine name with **CFA**, which replaces the parameter field address with the corresponding compilation address.

In **FINISH.HAIRCUT** we add one to **SERVED**. If a customer is waiting we decrement **WAITING** by one and schedule another haircut completion, otherwise we set the status of **BARBER** to idle and exit. Finally, we include the word **REPORT** to display the current status of the system—number of customers served, number waiting, and number of potential customers who have balked.

Instrumentation and Software

The example shown was run on a Commodore 64TM with a 1541 disk drive and a VIC-1525 graphic printer. The Forth system was Superforth 64TM V 2.2 from Parsec Research, run in its 79-Standard mode. Superforth 64 is an extension of MVP Forth from Mountainview Press. With the exceptions noted in the text, every effort was made to adhere to the Forth-79 standard. The screens shown here use the Brodie backslash (\, ASCII 92) to indicate that the rest of the line is comments [BRO84]. The actual screens used the Commodore code 92, which prints as an English pound sign (£).

Figure 3 shows some of the output from **LOADING** screen 97, including both computer system performance figures and simulation results. An additional 86 bytes not shown was consumed by the event calendar in the course of the simulation.

```
SYSTEM COMPILATION
727 BYTES 34.68 SECONDS

MODEL COMPILATION
601 BYTES 19.50 SECONDS

AFTER 8 HOURS 0 MINUTES 0 SECONDS.
26 CUSTOMERS HAVE BEEN SERVED,
0 ARE WAITING AND
0 HAVE BALKED.
9.03 SECONDS EXECUTION.

END OF SIMULATION.
```

Figure 3.

Conclusions

We have shown that a discrete event simulation system can be implemented in Forth using very modest amounts of memory and with acceptable run times. In addition, the nature of the Forth language makes the system inherently interactive and extendable, features not found in more conventional simulation systems.

References

- [BRO81] Brodie, Leo, *Starting Forth*, Prentice-Hall Inc., Englewood Cliffs, New Jersey, 1981.
- [BRO84] Brodie, Leo, *Thinking Forth, A Language and Philosophy for Solving Problems*, Prentice-Hall Inc., Englewood Cliffs, New Jersey, 1984.
- [KNU72] Knuth, Donald E., *The Art of Computer Programming Vol. 1, Fundamental Algorithms*, Addison-Wesley, Reading, Mass. 1972.
- [MCC81] McCormack, William M. and Robert G. Sargent, "Analysis of Future Event Set Algorithms for Discrete Event Simulation", *Communications of the ACM*, Vol. 24, #12:801-812, Dec. 1981.
- [PET81] Petersen, Joel V., "Recursion and the Ackermann Function", *Forth Dimensions*, Vol. III, #3:89-90, 1981.
- [PRI74] Pritsker, A. Alan B., *The GASPIV Simulation Language*, John Wiley & Sons, New York, NY 1974.
- [SCH74] Schriber, Thomas J., *Simulation Using GPSS*, John Wiley & Sons, New York, NY 1974.
- [SHE84] Sheppard, Sallie, Udo Pooch and C. Dennis Pegden, eds., *1984 Winter Simulation Conference Proceedings*, Society for Computer Simulation, La Jolla, Calif. 1984.
- [SIG84] *1984 ACM SIGMETRICS Conference on Measurement and Modeling of Computer Systems*. ACM Order #488840, ACM, Baltimore, Md.
- [SUR83] Suri, Rajan "Infinitesimal Perturbation Analysis of Discrete Event Dynamic Systems: A General Theory", *Proceedings of 22nd IEEE Conference on Decision and Control*. December 1983.

Mr. Zettel has a B.S. in Chemical Engineering in 1959 from a school he would just as soon not name. He also has an MBA (1975) from Canisius College. Mr. Zettel did his first simulation model in 1962, using GPSS to study the interaction of specification and aging constraints during propellant filling on the final quality of Polaris rocket motors while working for the Aerojet-General Corporation.

He has since done simulation studies for a number of companies using various simulation languages and techniques, including GPSS, GASP IV, SLAM II, and SIMAN. Currently a member of the research staff of the Ford Motor Company, Mr. Zettel's interests center on the simulation of manufacturing systems.

SCREEN #17

```

0 \ RANDOM NUMBER GENERATOR
1 \ FROM STARTING FORTH P265.
2 CREATE RND 0 , HERE RND !
3 : RANDOM RND @ 31421 * 6927 + DUP
4   RND ! ;
5 : CHOOSE ( U1 --- U2)
6   RANDOM U* SWAP DROP ;
7
8 : <ROT (N1 N2 N3...N3 N2 N1) \ BACKWARDS ROT
9   ROT ROT ;
10
11
12
13
14
15

```

SCREEN #36

```

0 \ FREE MIN.NODE GET
1 VARIABLE FREE 0 FREE ! 3 CONSTANT MIN.NODE
2 : GET ( N --- ADDR) \ PUT THE ADDRESS OF THE NEXT AVAILABLE
3 \ NODE OF SIZE N ON THE STACK.
4 DUP FREE DUP @ \ FIND A FREE NODE BIG ENOUGH.
5 BEGIN DUP
6   IF DUP 2+ @ 4 PICK <
7   ELSE DUP THEN
8 WHILE SWAP DROP DUP @ REPEAT ?DUP
9 IF DUP 2+ @ 4 PICK - DUP MIN.NODE < \ SALVAGE REST OF NODE?
10 IF DROP DUP @ ROT ! SWAP DROP
11 ELSE OVER 5 ROLL 2* + DUP 5 ROLL ! 3 PICK @ OVER ! 2+ !
12 THEN
13 ELSE DROP HERE SWAP 2* ALLOT \ NO NODE; MAKE ONE.
14 THEN DUP <ROT 2+ ! ;
15

```

SCREEN #37

```

0 \ N! N@ 2+! GIVE
1 : N! ( N1...NN ADDR N ---)
2 \ STORE N1 TO NN IN CONSECUTIVE CELLS STARTING AT ADDR.
3 2* 0 DO SWAP OVER I + ! 2 +LOOP DROP ;
4 : N@ ( ADDR N --- N1...NN) \ FETCH N CONSECUTIVE VALUES STARTING
5 \ AT ADDR + 2N-2 & LEAVE THEM ON THE STACK.
6 1- 2* 0 SWAP DO DUP I + @ SWAP -2 +LOOP DROP ;
7
8 : 2+! ( D1 ADDR ---) \ ADD D1 TO THE DOUBLE NUMBER AT ADDR.
9 DUP >R 2@ D+ R> 2! ;
10
11 : GIVE ( ADDR ---)
12 \ ADD NODE AT ADDR TO FREE LIST.
13 FREE @ OVER ! FREE ! ;
14
15

```

SCREEN #38

```

0 \ PUSH POP
1 : PUSH ( N1...NN ADDR ---) \ PUSH N1 TO NN ONTO THE STACK
2 \ POINTED TO BY ADDR. NN IS THE NODE SIZE IN CELLS.
3 OVER GET OVER @ SWAP DUP 4 ROLL !
4 3 PICK N! ;
5
6 : POP ( ADDR --- N1...NN) \ POP STACK POINTED TO BY ADDR,
7 \ LEAVING VALUES ON STACK & NODE ON FREE LIST.
8 DUP @ DUP NOT ABORT" EMPTY USER STACK."
9 DUP @ ROT ! DUP GIVE 2+ DUP @ 1- N@ ;
10 : ME ( --- ADDR) \ PUT THE CFA OF THE WORD BEING
11 \ COMPILED ON THE STACK.
12 LATEST PFA CFA [COMPILE] LITERAL ; IMMEDIATE
13
14
15

```

SCREEN #39

```

0 \ 2PICK CALENDAR SCHEDULE
1 : 2PICK ( N--- D1) \ RETURN THE CONTENTS OF THE NTH AND N+1ST
2 \ STACK VALUES, NOT COUNTING N.
3 1+ DUP 1+ PICK SWAP PICK ;
4
5 VARIABLE CALENDAR 0 CALENDAR !
6 : SCHEDULE ( N1...NN ADDR DTIME M ---) \ PUT THE EVENT WITH
7 \ COMPILATION ADDRESS ADDR ON THE EVENT CALENDAR.
8
9 CALENDAR DUP @
10 BEGIN DUP IF 4 2PICK 3 PICK 4 + 2@ D>
11 ELSE DUP THEN
12 WHILE SWAP DROP DUP @ REPEAT
13 DROP PUSH ;
14
15

```

SCREEN #40

```

0 \ CLOCK NEXT.EVENT SIMULATE END.RUN
1 2VARIABLE CLOCK
2 : NOW ( --- D1) \ D1 IS THE CURRENT SIMULATION CLOCK TIME.
3  CLOCK 2@ ;
4 : NEXT.EVENT ( ---) \ TAKE THE NEXT EVENT OFF THE CALENDAR,
5 \ UPDATE THE SIMULATION CLOCK, AND EXECUTE THE EVENT.
6  CALENDAR @ NOT ABORT" NO NEW EVENT"
7  CALENDAR POP DROP CLOCK 2! EXECUTE ;
8
9 : SIMULATE ( ---)
10 BEGIN NEXT.EVENT ?TERMINAL UNTIL ;
11
12 : END.RUN ( ---)
13  ABORT" END OF SIMULATION." ;
14
15

```

SCREEN #88

```

0 \ TRUE FALSE WAITING BARBER BALKS SERVED FINISH.HAIRCUT
1 1 CONSTANT TRUE 0 CONSTANT FALSE
2 VARIABLE WAITING VARIABLE BARBER
3 VARIABLE BALKS VARIABLE SERVED
4 VARIABLE ARRIVAL.MIN VARIABLE ARRIVAL.SPREAD
5 VARIABLE HAIRCUT.MIN VARIABLE HAIRCUT.SPREAD
6 : NOW. ( ---) \ PRINT THE CURRENT SIMULATION TIME.
7  NOW 3600 U\MOD
8  . ." HOURS " 60 \MOD . ." MINUTES " . ." SECONDS. " ;
9 : FINISH.HAIRCUT ( ---)
10 1 SERVED +! WAITING @
11 IF -1 WAITING +! ME HAIRCUT.MIN @ HAIRCUT.SPREAD @ CHOOSE
12   + S->D NOW D+ 5 SCHEDULE
13 ELSE FALSE BARBER !
14 THEN ;
15

```

SCREEN #89

```

0 \ CUSTOMER.ARRIVES SETUP.BARBERSHOP
1 : CUSTOMER.ARRIVES ( ---)
2  ME ARRIVAL.MIN @ ARRIVAL.SPREAD @ CHOOSE +
3  S->D NOW D+ 5 SCHEDULE BARBER @
4  IF WAITING @ 4 <
5    IF 1 WAITING +! ELSE 1 BALKS +! THEN
6  ELSE TRUE BARBER ! ["] FINISH.HAIRCUT CFA HAIRCUT.MIN @
7    HAIRCUT.SPREAD @ CHOOSE + S->D NOW D+ 5 SCHEDULE
8  THEN ;
9 : SETUP.BARBERSHOP ( ---) \ INITIALIZE THE SIMULATION.
10 0 WAITING ! FALSE BARBER ! 0 BALKS ! 0 SERVED !
11 0. CLOCK 2! CUSTOMER.ARRIVES ;
12
13
14
15

```


SCREEN #96

```
0 \ BARBERSHOP REPORT
1 : REPORT ( ---)
2 CR ." AFTER " NOW. CR SERVED @ .
3 ." CUSTOMERS HAVE BEEN SERVED," CR
4 WAITNG @ . ." ARE WAITING AND" CR BALKS @ .
5 ." HAVE BALKED." CR ;
6
7
8
9
10
11
12
13
14
15
```

SCREEN #97

```
0 \ DEMO LOAD SCREEN FOR BARBER SHOP SIMULATION.
1 : ELAPSED.TIME ( D1 D2 ---) 2SWAP D- DROP 100 60 */ S->D
2 <# # # 46 HOLD #S #> TYPE ." SECONDS " ;
3 HERE RDTIM 17 DUP . LOAD 36 40 THRU
4 CR ." SYSTEM COMPILATION" CR ROT HERE SWAP - . ." BYTES "
5 RDTIM ELAPSED.TIME HERE RDTIM
6 88 89 THRU 96 DUP . LOAD
7 CR ." MODEL COMPILATION" CR ROT HERE SWAP - . ." BYTES " RDTIM
8 ELAPSED.TIME RDTIM
9 : END.SIM ( D1 ---) REPORT RDTIM ELAPSED.TIME ." EXECUTION." CR
10 ABORT" END OF SIMULATION" ;
11 720 ARRIVAL.MIN ! 720 ARRIVAL.SPREAD !
12 720 HAIRCUT.MIN ! 480 HAIRCUT.SPREAD !
13 SETUP.BARBERSHOP
14 ' END.SIM CFA 28800. 5 SCHEDULE SIMULATE
15
```

