STAN-CS-73-352 SU-SEL-73-011

The Expected Difference Between the SLTF and MTPT Drum Scheduling Disciplines

(Digest Edition)

bY

Samuel H. Fuller

August 1972

Technical Report No. 28

This work was conducted while the author was a Hertz Foundation Graduate Fellow at Stanford University, Stanford, Calif., and was partially supported by the Joint Services Electronics Program, U.S. Army, U.S. Navy, and U.S. Air Force under Contract N-00014-67-A-0112-0044. Computer time was provided by the Atomic Energy Commission under contract AT-(0403) 515.

DIGITAL SYSTEMS LABORATORY

STANFORD ELECTRONICS LABORATORIES

STANFORD UNIVERSITY . STANFORD, CALIFORNIA

STAN-CS-73-352 SEL 73-011

THE EXPECTED DIFFERENCE BETWEEN SLTF AND MTPT DRUM SCHEDULING DISCIPLINES

(Digest edition)

by

Samuel H. Fuller

August 1972

Technical Report no. 28

DIGITAL SYSTEMS LABORATORY

Dept. of Electrical Engineering Dept. of Computer Science

Stanford University

Stanford, California

This work was conducted while the author was a Hertz Foundation Graduate Fellow at Stanford University, Stanford, Calif., and was partially supported by the Joint Services Electronics Program, U.S. Army, U. S. Navy, and U.S. Air Force under Contract N-00014-67-A-0112-0044. Computer time was provided by the Atomic Energy Commission under contract AT-(O&-3)515.

THE EXPECTED DIFFERENCE BETWEEN THE SLTF AND MTPT

DRUM SCHEDULING DISCIPLINES

ABSTRACT

This report is a sequel to an earlier report [Fuller, 1971] that develops a minimal-total-processing-time (MTPT) drum scheduling algorithm. A quantative comparison between MTPT schedules and shortest-latency-time-first (SLTF) schedules, commonly acknowledged as good schedules for drum-like storage units, is presented here. The analysis develops an analogy to random walks and proves several asymptotic properties of collections of records on drums. These properties are specialized to the MTPT and SLTF algorithms and it is shown that for sufficiently large sets of records, the expected processing time of a SLTF schedule is longer than a MTPT schedule by the expected record length. The results of a simulation study are also presented to show the difference in MTPT and SLTF schedules for small sets of records and for situations not covered in the analytic discussion.

DISCUSSION

In Fuller [1971] we introduced a drum scheduling algorithm that can efficiently find schedules for sets of I/O requests that minimize the total rotational delay (latency) of the set of I/O requests. The original article, however, is entirely devoted to developing the scheduling algorithm, proving its correctness, and presenting a few examples of the algorithm in operation; this article provides a quantative measure of how much better the new drum scheduling algorithm can be expected to be over conventional scheduling algorithms.

First, briefly reconsider the scheduling problem posed in the original paper. Suppose a fixed-head drum, as illustrated in Fig. 1.1, receives requests to process N I/O records. These requests may be to either read or write a record onto the drum; no distinction is made between reading or writing in this, or the original, discussion. In Fig. 1.1, notice we allow the records to start anywhere around the circumference of the drum and furthermore the record lengths are arbitrary. We assume the drum can only begin reading a record at s_i, the record's starting address, and once started, the drum cannot be pre-empted and will finish processing the record at f_i, the finishing address. The interval of time the drum is delayed waiting for the beginning of the next record is called <u>rotational latency</u> or simply latency. Furthermore, we exclude the possibility of more I/O requests arriving at the drum while the original N requests are being serviced. This is an unrealistic assumption in some cases and more will be said about this in Sec. 6, but for the present we will forbid random arrivals. A scheduling algorithm is developed in the original paper that finds a schedule that processes all N records in the minimal amount of time, and hence we will denote such

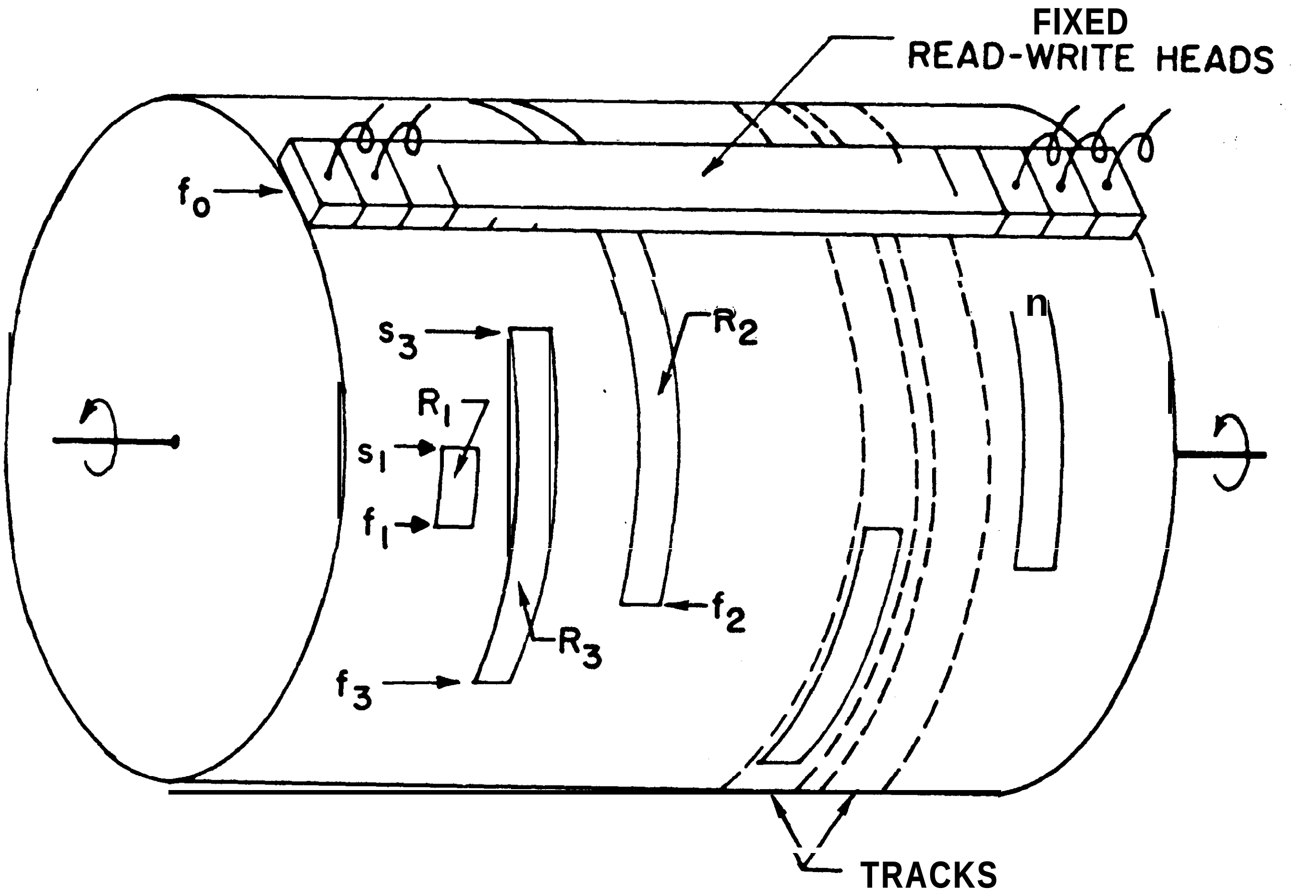

Figure 1.1. A drum storage unit.

a schedule as a minimal-total-processing-time (MTPT) schedule.* An important property of this scheduling algorithm is that it is able to construct a MTPT schedule in N*logN simple steps.

The algorithm that is commonly acknowledged as a good drum scheduling algorithm is shortest-latency-time-first (SLTF); for this reason the MTPT algorithm will be compared to the SLTF algorithm in this article. A SLTF schedule is simply a schedule that processes the next record to come under the read-write heads, given that the read-write heads are not busy servicing another request. In general, an SLTF schedule is not a MTPT schedule and this article investigates how much longer than a MTPT schedule a SLTF schedule needs to process a set of N records. Specifically, this article presents an asymptotic expression for the expected difference between the SLTF and MTPT schedules and then concludes with some empirical results to show how the expected difference of the two algorithms behaves before it approaches its asymptotic value. These results, along with the least upper bound of one drum revolution for the difference between SLTF and MTPT schedules developed earlier [Stone and Fuller, 1971], places us in a good position to quantatively evaluate the relative advantages offered by either the MTPT or SLTF schedules when minimizing the total processing time of a set of I/O requests is a reasonable objective.

---

* The algorithm was called an optimal drum scheduling algorithm in the original article, but this article refers to the algorithm as the minimal-total-processing-time (MTPT) drum scheduling algorithm. This name is more mnemonic and recognizes that other drum scheduling algorithms may be optimal for other optimality criteria.

REFERENCES

Fuller, S. H. (1971) An optimal drum scheduling algorithm. Technical Report no. 12, Digital Systems Laboratory, Stanford University, Stanford, California (April 1971).

Stone, H. S. and S. H. Fuller (1971) On the near-optimality of the shortest-access-time-first drum scheduling discipline. Technical Note no. 12, Digital Systems Laboratory, Stanford University, Stanford, California (October 1971).