May, 1984 LIDS-P-1372

INDIVIDUALLY OPTIMAL ROUTING IN PARALLEL SYSTEMS***

P.R. KUMAR* AND J. WALRAND**

ABSTRACT

Jobs arrive at a buffer from which there are several parallel routes to a destination. A socially optimal policy is one which minimizes the average delay of all jobs, whereas an individually optimal policy is one, which for each job, minimizes its own delay, with route preference given to jobs at the head of the buffer. If there is a socially optimal policy for a system with no arrivals, which can be implemented by each job following a policy γ in such a way that no job ever utilizes a previously declined route, then we show that such a γ is an individually optimal policy for each job. Moreover γ continues to be individually optimal even if the system has an arbitrary arrival process, subject only to the restriction that past arrivals are independent of future route traversal times. Thus, γ is an individually optimal policy which is insensitive to the nature of the arrival process. In the particular case where the times to traverse the routes are exponentially distributed with a possibly different mean time for each of the parallel routes, then such an insensitive individually optimal policy does in fact exist and is moreover trivially determined by certain threshold numbers. A conjecture is also made about more general situations where such individually optimal policies exist.

* Dept. of Mathematics and Computer Science, University of Maryland, Baltimore County, Baltimore, Maryland 21228, U.S.A. Currently visiting L.I.D.S., M.I.T., Cambridge, MA 02139

** Dept. of EECS, Univ. of California, Berkeley, CA 94720.

*** The research of the first author has been supported by the U.S.A.R.O. under Contract No. DAAG-29-84-K0005.

Submitted to SIAM Journal on Control and Optimization

I. INTRODUCTION

We consider the system of Figure 1. Jobs arrive into

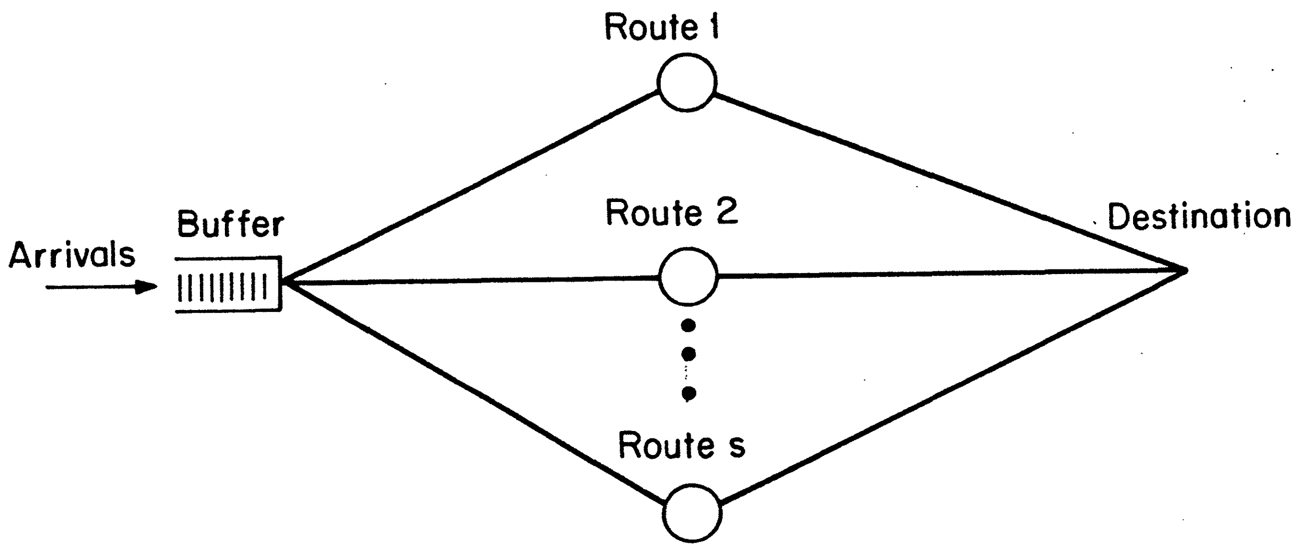

FIGURE 1

a buffer from which there are s parallel routes to a destination. Each job wishes to minimizes its own delay, which is the expected duration of time from the arrival of the job to the time at which it reaches its destination. Every route can accomodate only one job at a time, and a route is idle if no job is currently on that route.

At every time instant, all the currently idle routes are offered to the job occupying the first position, and it can either choose to traverse one of them or decline all the currently idle routes and continue to wait. Then, all the (remaining) idle routes are offered to the job in the second position, then the third position and so on till the end of the buffer.

A policy for a job which specifies when it should choose or decline an offered idle route, and which minimizes the delay for that job, will be called an <u>individually optimal policy</u>. We seek to determine such individually optimal policies.

Instead of considering such individually optimal policies, one can also seek to obtain a policy which minimizes the average delay of <u>all</u> jobs. Such a policy will be called a <u>socially optimal policy</u>.

In general, the notions of social and individual optimality do not coincide, because a socially optimal policy may "sacrifice" one job if by doing so it can lessen the overall delays of all jobs. In contrast, a job implementing an individually optimal policy will never make such sacrifices.

Our main results are the following.

<u>Theorem 1</u>

Assume that

(1.i) The system has no arrivals.

(1.ii) There is a socially optimal policy π which can be implemented by each job individually implementing a policy γ.

(1.iii) Under γ, if a job ever declines an offered idle route, then it will never thereafter utilize that route.

Then, γ is an individually optimal policy for this system.

<u>Theorem 2</u>

The policy γ of Theorem 1 is not only optimal for a system with no arrivals, but is also optimal for all arrival processes which satisfy:

Past arrivals are independent of future route traversal times. (2)

Theorem 3

Suppose that

(3.i) For $i=1,2,\ldots,s$ the time to traverse route i is exponentially distributed with mean μ_i^{-1} . (3)

(3.ii) (Without loss of generality) $\mu_1 \geq \mu_2 \geq \ldots \geq \mu_s$.

and the arrival process satisfies (2). Define

$$T_j := \frac{\mu_1+\mu_2+\ldots+\mu_{j-1}}{\mu_j} - (j-1) \tag{4}$$

Then, an individually optimal policy γ is given by the following rule:

Let j be the fastest idle route offered to a job which is in position c in the buffer. Then the job should utilize route j if and only if $c > T_j$

It can be noted that by identifying each route i of Figure 1 with a server, Theorems 1 and 2 are really studying the optimal utilization, individually by each customer or socially, of GI/G/S queues where each of the servers is allowed to be different. By the same analogy, Theorem 3 gives the optimal policy for utilization by individual customers of GI/M/S queues with heterogenous servers.

Recently, Agrawala, Coffman, Garey and Tripathi [1] considered the problem of obtaining a socially optimal policy when there are no arrivals to the buffer, and the times to traverse the routes satisfy

(3.i,ii). They have shown that the socially optimal policy π which minimizes the total delay of all existing jobs is given by the following rule:

> Let j be the fastest idle route available when c customers are waiting in the buffer. Then assign one customer from the buffer to route j if and only if $c>T_j$. (6)

It can be seen that the socially optimal policy (6) when there are no arrivals, and the individually optimal policy (5) when arrivals according to an arbitrary process satisfying (2) are allowed, are closely linked by Theorems 1 and 2.

Progress has also recently been made on the problem of obtaining socially optimal policies for exponential routes when the arrivals form a Poisson process. Larsen [2] conjectured that if the arrivals are Poisson and (3) is satisfied, then the socially optimal policy would be of threshold type. For the case s=2 where there are two exponential routes with means $\mu_1^{-1} \leq \mu_2^{-1}$, this conjecture has been proved by Lin and Kumar [3], i.e. there is a number τ such that the second route should be utilized if and only if the number of jobs in the buffer exceeds τ. It has been conjectured that as the rate λ of the Poisson arrival process decreases, τ increases and converges to T_2 as $\lambda\to 0$, see [3], but this result has still not been rigorously proved. Recently, Walrand [4] has obtained another proof of the threshold result of [3] using stochastic coupling arguments instead of the policy iteration argument used in [3]. For the case of s>2 routes, the problem is still open. In Lin and Kumar [3] it is conjectured that the socially optimal policy is given by threshold functions $\tau_2(\cdot)$, $\tau_3(\cdot),\ldots,\tau_s(\cdot)$ such that route j should be utilized if and only if the number c of customers in the buffer satisfies $c>\tau_j(\vec{b})$ where $\vec{b}$ is the vector describing which routes are idle and which are busy. This also is unproved to date. Lastly, it again appears that $\tau_j(\vec{b})$ increases as $\lambda \searrow 0$ and converges to T_j.

II. A FALLACIOUS ARGUMENT

Before presenting the proof, we present a tempting argument just for Theorem 3, which suggests itself, but which is fallacious.

Suppose that a job is in position c and route j is the fastest idle route offered to it. If this job chooses route j, then its delay is μ_j^{-1}. Suppose however that the job declines route j and decides instead to wait for a faster route to become available. Since there are c-1 jobs ahead of the job being considered, it is tempting to analyze what is the delay of the job if it "chooses" to wait till (c-1) completions have occured from routes 1,2,...,j-1 and then utilize the next subsequent route which becomes available. The expected time to wait till c=(c-1)+1 completions have occured from routes 1,2,...,j-1 is

$$\frac{c}{\mu_1+\mu_2+\dots+\mu_{j-1}}$$

Now, the probability that route i is the one which becomes available at the c-th completion is

$$\frac{\mu_i}{\mu_1+\mu_2+\dots+\mu_{j-1}}$$

and if it chooses this route, its expected route traversal time is

$$\frac{1}{\mu_i}$$

Hence, if the job can indeed follow this policy, then its delay is

$$\frac{c}{\mu_1+\ldots+\mu_{j-1}} + \sum_{i=1}^{j-1} \frac{\mu_i}{\mu_1+\ldots+\mu_{j-1}} \cdot \frac{1}{\mu_i} = \frac{c+j-1}{\mu_1+\ldots+\mu_{j-1}}$$

Thus, comparing the two alternatives of choosing route j now, or waiting till c completions have occurred from routes 1,2,...,j-1 and choosing whichever route is the c-th completion, we obtain

$$\text{if } \frac{c+j-1}{\mu_1+\mu_2+\ldots+\mu_{j-1}} > \frac{1}{\mu_j}, \quad \text{then choose route } j$$

$$\leq \frac{1}{\mu_j}, \quad \text{then decline route } j$$

which turns out to be the same comparison as

$$\text{if } c > T_j \qquad \text{then choose route } j$$

$$\leq T_j \qquad \text{then decline route } j$$

suggested in Theorem 3.

Unfortunately, the above reasoning is fallacious. Firstly, the job may not even be allowed to use the route which becomes available at the c-th completion from routes 1,2,...,j-1. The reason is that there may be other jobs which are ahead of it in the buffer at the time of the c-th completion, which occurs if other jobs ahead of it have declined to choose some of the routes which have become available in the interim, and one of these jobs ahead of it may decide to use this route. Secondly, if the c-th completion corresponds to a "slow route", then the job by this time may have advanced so far to the front of the buffer, that it may prefer to decline the c-th completion.

It is therefore somewhat surprising that the policy of Theorem 3 is indeed individually optimal.

III. PROOFS

We shall first prove Theorems 1 and 2 simultaneously. The proof is by induction. Suppose that there is only one job in the buffer and there are no arrivals to the buffer. Then individual and social optimality coincide, and so by (1.ii), the policy γ is individually optimal in this situation.

One property is important.

(P) Since there is no one in the buffer behind the job, and no future arrivals, the job can always utilize a route that it previously chose to decline. This is because an idle route remains idle forever if this one job has declined it.

Suppose now that there are possibly other jobs in the buffer behind the job in position 1, and arrivals satisfying (2) can occur. Then, due to interference by other jobs, the option outlined in (P) ceases to be available to the job initially in position 1. Since the set of options available has decreased, the delay of the job in position 1 can only increase over what it was in the situation where it was the only job in the buffer and no arrivals could occur. However, by (1.iii), the options eliminated are not really needed when γ is used, and so the delay remains the same as what it was. Hence, the suggested policy γ is optimal for the job in position 1, irrespective of the total number of jobs in the buffer and irrespective of the arrival process, so long as it satisfies (2).

Suppose, for purposes of induction, that we have already shown that the suggested policy is optimal for positions 1,2,...,k-1 regardless of the total number of jobs in the buffer and the arrival process satisfying (2).

Now consider a job in position k, and suppose that the total number of jobs in the buffer is also k, and no arrivals are allowed. Jobs 1,2,...,k-1 behave according to their individually optimal policies and, by property (1.iii), their behaviour is unaffected by what the job in position k does. If k had a strictly better policy than the one suggested, then its delay would strictly decrease, while the delays of the jobs in positions 1,2,...,k-1 is unchanged. But this is impossible, because then the total delay of all k jobs would be strictly less than what it would be under the policy already known to be socially optimal in this situation. Hence the suggested individual policy is optimal for the job in position k, when there are no other jobs behind it in the buffer, and there are no arrivals to the buffer.

Now suppose that there are possibly other jobs in the buffer behind the job in position k and also arrivals are allowed as in (2). Since (P) does not hold anymore for the job currently in position k, its delay can only increase in comparison with what it was when it did not have any other jobs behind it in the buffer and no arrivals were allowed. However, by implementing the suggested policy γ, (1.iii) ensures that the eliminated options are unnecessary anyway, and so the job in position k can ensure that its delay remains unchanged in comparison with the situation where there were no other jobs behind it in the buffer, and no arrivals were allowed. Hence the suggested policy γ is optimal for the job in position k, irrespective of the total number of jobs in the buffer, and irrespective of the arrival process, so long as it satisfies (2).

The induction is now complete, and we have proved that the suggested policy γ is individually optimal in the situations of Theorems 1 and 2.

Now we turn to the proof of Theorem 3. We use the result of Agrawala, Coffman, Garey and Tripathi [1] that the policy π of (6) is socially optimal when (3.i,ii) are true and no arrivals are present. Now consider the policy γ suggested in (5). First, we show that when all jobs implement γ, then π will be implemented, which will show that (1.ii) is satisfied. Suppose there are c jobs in the buffer and route j is the fastest idle route available. If $c \leq T_j$ then all jobs in the buffer will, one by one, decline route j since they are following the rule (5) and their positions are all less than or equal to c, and so (6) is implemented. If $c > T_j$, then one job, namely the ℓ-th one where ℓ is the smallest integer greater than T_j will utilize route j, and so again (6) is implemented. Hence (1.ii) is satisfied. Now we show (1.iii). Let k be the position of job which declines route j, the fastest idle route offered to it. Since γ is implementing (5), $k < T_j$. As time goes on, the job, if it has not already utilized some route, can only improve its position, i.e. its position can only decrease. Since T_j remains fixed, its position will never thereafter exceed T_j and so route j will never be utilized by this job. Hence (1.iii) is satisfied, and the proof Theorem 3 is complete.

IV. CONCLUDING REMARKS

We do not expect that the condition (1.iii) will hold in all situations. We have however seen one situation, namely the exponential routes case of Theorem 3 where conditions (1.ii,iii) are satisfied.

Are there other situations of more general "service" times, where conditions (1.ii,iii) hold? The problem here is that not much progress has been made on the problem of social optimality when there are no arrivals, other than in the exponential routes case of [1].

We conjecture that if

(7.i) the hazard rate of each server is an increasing function

(7.ii) The range $[\underline{h}_i, \bar{h}_i]$ of the hazard rate of each server i is such that $\bar{h}_{i+1} \leq \underline{h}_i$, (7)

then, the socially optimal policy in the case of no arrivals will indeed satisfy (1.ii,iii). Thus, we conjecture that the strong result of Theorem 2, namely the existence of an individually optimal policy insensitive to the arrival process, will be true when (7.i,ii) are satisfied. It is also reasonable to guess that it will be of threshold type, where the threshold function depends on the server states.

REFERENCES

[1] A.K. Agrawala, E.G. Coffman, Jr., M.R. Garey and S.K. Tripathi, "A Stochastic Optimization Algorithm Minimizing Expected Flow Times on Uniform Processors", Preprint, 1983.

[2] R.L. Larsen, "Control of Multiple Exponential Servers with Application to Computer Systems", Ph.D. Thesis, Technical Report No. TR-1041, University of Maryland, April 1981.

[3] W. Lin and P.R. Kuman, "Optimal Control of a Queueing System with Two Heterogeneous Servers," to appear in IEEE Transactions on Automatic Control, 1984.

[4] J. Walrand, "A Note On "Optimal Control of a Queueing System with Two Heterogeneous Servers," to appear in Systems and Control Letters.