

On the rod placement theorem of Rybko and Shlosman

Venkat Anantharam

Received: 18 May 2005 / Accepted: 14 September 2005
© Springer Science + Business Media, LLC 2006

Abstract Given $n - 1$ points $x_1 \leq x_2 \leq \dots \leq x_{n-1}$ on the real line and a set of n rods of strictly positive lengths $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$, we get to choose an n -th point x_n anywhere on the real line and to assign the rods to the points according to an arbitrary permutation π . The rod $\lambda_{\pi(k)}$ is thought of as the workload brought in by a customer arriving at time x_k into a first in -first out queue which starts empty at $-\infty$. If any x_i equals x_j for $i < j$, service is provided to the rod assigned to x_i before the rod assigned to x_j .

Let $Y_\pi(x_n)$ denote the set of departure times of the customers (rods). Let $N_\pi(x_1, \dots, x_{n-1}; \lambda_1, \dots, \lambda_n)$ denote the number of choices for the location of x_n for which $0 \in Y_\pi(x_n)$. Rybko and Shlosman proved that

$$\sum_{\pi} N_\pi(x_1, \dots, x_{n-1}; \lambda_1, \dots, \lambda_n) = n!$$

for Lebesgue almost all $(x_1, \dots, x_{n-1}; \lambda_1, \dots, \lambda_n)$.

Let $y_{\pi,k}(x_n)$ denote the departure point of the rod λ_k . Let $N_{\pi,k}(y)$ denote the number of choices for the location of x_n for which $y_{\pi,k}(x_n) = y$ and let $N_k(y) = \sum_{\pi} N_{\pi,k}(y)$. In this paper we prove that for every $(x_1, \dots, x_{n-1}; \lambda_1, \dots, \lambda_n)$ and every k we have $N_k(y) = (n - 1)!$ for all but finitely many y . This implies (and strengthens) the rod placement theorem of Rybko and Shlosman.

Keywords FCFS queue · FIFO queue · Rod placement · Symmetric group

Research supported by ONR MURI N00014-1-0637, NSF ECS-0123512, Marvell Semiconductor, and the University of California MICRO program.

V. Anantharam (✉)
EECS Department, University of California, Berkeley, CA 94720,
USA
e-mail: ananth@eecs.berkeley.edu

AMS Subject Classifications: 60G55, 05A05, 60C05, 60K25

Discussion

In this paper we prove a result which implies (and strengthens) the rod placement theorem of Rybko and Shlosman ([2]). For the basic facts we use from the theory of queues see, for instance, Asmussen ([1]).

Theorem . Given $n - 1$ points $x_1 \leq x_2 \leq \dots \leq x_{n-1}$ on the real line and a set of n rods of strictly positive lengths $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$, we get to choose an n -th point x_n anywhere on the real line and to assign the rods to the points according to an arbitrary permutation π . The rod $\lambda_{\pi(k)}$ is thought of as the workload brought in by a customer arriving at time x_k into a first in -first out (FIFO) queue which starts empty at $-\infty$. If any x_i equals x_j for $i < j$, service is provided to the rod assigned to x_i before the rod assigned to x_j .

Let $y_{\pi,k}(x_n)$ denote the departure point of the rod λ_k . Let $N_{\pi,k}(y)$ denote the number of choices for the location of x_n for which $y_{\pi,k}(x_n) = y$ and let $N_k(y) = \sum_{\pi} N_{\pi,k}(y)$. Then for every $(x_1, \dots, x_{n-1}; \lambda_1, \dots, \lambda_n)$ and every k we have $N_k(y) = (n - 1)!$ for all but finitely many y .

Remark 1. Let $N_\pi(x_1, \dots, x_{n-1}; \lambda_1, \dots, \lambda_n)$ denote the number of choices for the location of x_n for which $0 \in Y_\pi(x_n) = \{y_{\pi,1}(x_n), \dots, y_{\pi,n}(x_n)\}$. Rybko and Shlosman ([2]) proved that

$$\sum_{\pi} N_\pi(x_1, \dots, x_{n-1}; \lambda_1, \dots, \lambda_n) = n!$$

for Lebesgue almost all $(x_1, \dots, x_{n-1}; \lambda_1, \dots, \lambda_n)$. To derive this from the theorem, since $0 \notin B$ for Lebesgue almost all $(x_1, \dots, x_{n-1}; \lambda_1, \dots, \lambda_n)$, we have

$$N_\pi(x_1, \dots, x_{n-1}; \lambda_1, \dots, \lambda_n) = \sum_{k=1}^n N_{\pi,k}(0),$$

and so

$$\sum_\pi N_\pi(x_1, \dots, x_{n-1}; \lambda_1, \dots, \lambda_n) = \sum_{k=1}^n N_k(0) = n!.$$

Proof of the theorem. Call $(x_1, x_2, \dots, x_{n-1})$ the rigid points and x_n the free point. Define a finite set $B = B(x_1, \dots, x_{n-1}; \lambda_1, \dots, \lambda_n)$ of “bad” points as follows : for a rigid point x_i and a subset $\Sigma \subseteq \{1, \dots, n\}$, consider the point $x_i + \sum_{k \in \Sigma} \lambda_k$; B is comprised of all such points, as i ranges over $1 \leq i \leq n - 1$ and Σ ranges over the subsets of $\{1, \dots, n\}$.¹ The theorem will be shown to apply to all $y \notin B$.

We prove the theorem by induction on n . Start with $n = 1$. There are no rigid points, so $B = B(; \lambda_1)$ is empty. For every y there is exactly one choice for the free point, namely $y - \lambda_1$, such that the rod λ_1 , which must necessarily be assigned to the free point, departs at y .

To propagate the induction, suppose the theorem has been proved for all configurations of $m - 1$ rigid points and m positive rod lengths for all $1 \leq m \leq n - 1$. We will now prove the theorem for $(x_1, \dots, x_{n-1}; \lambda_1, \dots, \lambda_n)$.

Let $y \notin B$. We first claim that it suffices to consider $y > x_{n-1}$. Indeed, if $y < x_1$ then $y_{\pi,k} = y$ iff we set $x_n = y - \lambda_k$ and the permutation π assigns rod λ_k to x_n and the other $n - 1$ rods to the rigid points in some order, and there are exactly $(n - 1)!$ ways of doing this. Also, if $x_m < y < x_{m+1}$ for some $1 \leq m \leq n - 2$, for each of the $(n - 1)(n - 2) \dots (m + 1)$ ways of assigning $n - m - 1$ of the rods other than λ_k to the rigid points $x_{m+1} \leq \dots \leq x_{n-1}$, by inductive hypothesis there are exactly $m!$ choices for the location of the free point and assignments of the remaining rods to the remaining points (the free point and the remaining rigid points) which result in rod λ_k departing at time y ,² so the claim is proved.

Consider now $y \notin B$ with $y > x_{n-1}$. For $l \neq k$, consider the configuration $(x_1, \dots, x_{n-2}; \lambda_1, \dots, \lambda_{l-1}, \lambda_{l+1}, \dots, \lambda_n)$. By inductive hypothesis there are exactly $(n - 2)!$ choices for the location of the free point, which we will denote x_F , and for the assignment of rods other than λ_l to x_F and the

points $x_1 \leq \dots \leq x_{n-2}$, which assignment we will denote by η , such that rod λ_k departs at y .

For such x_F and η , suppose the rod λ_k was assigned to one of the rigid points x_i , $1 \leq i \leq n - 2$, under η . Then the busy cycle of the FIFO queue that was in effect at time x_i continues to be in effect throughout the interval (x_i, y) , which includes x_{n-1} . Thus, in the original problem $(x_1, \dots, x_{n-1}; \lambda_1, \dots, \lambda_n)$, choosing $x_n = x_F$ and assigning rods to points according to π which extends η by assigning rod λ_l to x_{n-1} , will also have rod λ_k departing at time y .

Next, for such x_F and η , suppose the rod λ_k was assigned to x_F and had to wait before beginning service. Then there must be some x_i , $1 \leq i \leq n - 2$, which initiated the busy cycle in which the rod λ_k , which arrives at x_F , begins waiting. This would then mean that $y \in B$, which we have explicitly disallowed. Thus this case cannot occur.

It remains to consider those choices of x_F and η for which the rod λ_k was assigned to x_F and began to be served immediately on arrival at time x_F . This means $x_F = y - \lambda_k$.

In such a case, suppose $x_F < x_{n-1}$. Then x_{n-1} arrives during the service period $(x_F, x_F + \lambda_k = y)$ of rod λ_k . Once again, in the original problem $(x_1, \dots, x_{n-1}; \lambda_1, \dots, \lambda_n)$, choosing $x_n = x_F$ and assigning rods to points according to π which extends η by assigning rod λ_l to x_{n-1} , will have rod λ_k departing at time y .

Next, suppose $x_F > x_{n-1}$. Let z denote the end of the most recent busy cycle before x_F . If $(z \vee x_{n-1}) + \lambda_l < x_F$ the choice, in the original problem, of $x_n = x_F$ and the assignment of rods to points by π which extends η by assigning rod λ_l to x_{n-1} will again have rod λ_k departing at time y , because the service of rod λ_l will begin at $z \vee x_{n-1}$ and be completed before x_F . Note that this case can occur only if x_F and η are such that

$$y > (z \vee x_{n-1}) + \lambda_l + \lambda_k.$$

We are finally left with one case for x_F and η : the rod λ_k was assigned to x_F under η , began service immediately on arrival, and

$$(z \vee x_{n-1}) < x_F < (z \vee x_{n-1}) + \lambda_l,$$

where z denotes the end of the most recent busy cycle before x_F . Note that $x_F = y - \lambda_k$ in this case. Also note that this case can occur only if x_F and η are such that

$$y < (z \vee x_{n-1}) + \lambda_l + \lambda_k.$$

Thus, for any y either this case or the preceding case occurs, but not both.

In this case, in the original problem, we assign the rods to points according to π which equals η for all rods other than λ_k and λ_l , and assigns λ_k to x_{n-1} and λ_l to the free point.

¹ While it is not really necessary to include $\Sigma = \emptyset$ in the definition of B , it is convenient to not allow y in the theorem to be one of the rigid points.

² The bad set for the configuration of the remaining rigid points and the remaining rod lengths is contained in the original bad set B .

Note that

$$(z \vee x_{n-1}) + \lambda_k < y \quad \text{and} \\ y - [(z \vee x_{n-1}) + \lambda_k] < \lambda_l.$$

We place the free point, in the original problem, at the point $x_n < x_{n-1}$ such that

$$I(x_n, x_{n-1}) + y - [(z \vee x_{n-1}) + \lambda_k] = \lambda_l,$$

where $I(x_n, x_{n-1})$ denotes the total idle time of the FIFO queue associated to x_F and η over the interval (x_n, x_{n-1}) . With this choice of free point x_n and assignment π in the original problem, the rod λ_k departs at y .

We have now shown that, in the original problem, for each $l \neq k$ there are *at least* $(n - 2)!$ choices for x_n and an assignment π of rods to points such that rod λ_k departs at time y and either (1-1) x_{n-1} is assigned rod λ_l or (2-1) x_{n-1} is assigned rod λ_k and x_n is assigned rod λ_l . We will now prove that for each $l \neq k$ there are *exactly* $(n - 2)!$ choices for x_n and an assignment π of rods to points such that “(1-1) or (2-1) and rod λ_k departs at time y ” holds. Since the assignments satisfying “(1-1) or (2-1)” are disjoint as l ranges over $l \neq k$ and cover all possible assignments, this would complete the inductive step and the proof of the theorem.

Fix $l \neq k$ and a choice of x_n and an assignment π of rods to points such that “(1-1) or (2-1) and rod λ_k departs at time y ” holds. We show that each such x_n and π corresponds to a distinct choice of free point, which we denote x_F , and assignment, which we denote η , in the configuration $(x_1, \dots, x_{n-2}; \lambda_1, \dots, \lambda_{l-1}, \lambda_{l+1}, \dots, \lambda_n)$, which we call the *reduced* configuration, such that rod λ_k departs at time y in the reduced configuration. By the inductive hypothesis it would follow that there are at most $(n - 2)!$ such assignments, which would complete the proof of the theorem.

Suppose x_n and π are such that (1-1) holds and rod λ_k departs at time y in the original configuration. If π assigns rod λ_k to one of the original rigid points $x_i, 1 \leq i \leq n - 2$, the busy cycle containing x_i must last through the interval (x_i, y) , which contains x_{n-1} , so in the reduced configuration we could take $x_F = x_n$ and η to be the restriction of π that ignores the assignment of rod λ_l to x_{n-1} , and we would then have rod λ_k departing at time y in the reduced configuration.

Suppose x_n and π are such that (2-1) holds and rod λ_k departs at time y in the original configuration. If π assigns rod λ_k to x_n , it must be the case that rod λ_k does not have to wait before beginning service. Indeed, if it did have to wait before beginning service, this would have been because it arrived in some busy cycle which was initiated by some one of the rigid points $x_i, 1 \leq i \leq n - 1$, but that would then mean that $y \in B$, which we have explicitly disallowed.

Thus, we have $x_n = y - \lambda_k$. This then means that in the reduced configuration we could take $x_F = x_n$ and η to be the restriction of π that ignores the assignment of rod λ_l to x_{n-1} , and we would then have rod λ_k departing at time y in the reduced configuration.

Note that there are two ways in which the case just discussed could have occurred in the original configuration: either (a) $x_n < x_{n-1}$ or (b) $x_n > x_{n-1}$. In the former case our choice of x_F and η results in $x_F < x_{n-1}$. In the latter case, the fact that the rod λ_k , which was assigned to x_n , does not have to wait before beginning service means that

$$x_n > x_{n-1} + \lambda_l$$

and also

$$x_n > z + \lambda_l$$

where z denotes the departure time of the rod assigned to x_{n-2} . Recall that $x_n = y - \lambda_k$. Thus, this case can only occur if we had

$$y > (z \vee x_{n-1}) + \lambda_k + \lambda_l.$$

Also note that z can be computed based purely on the restriction of the assignment π (or equivalently η) to the rigid points $x_i, 1 \leq i \leq n - 2$.

Finally, suppose x_n and π are such that (2-1) holds and rod λ_k departs at time y in the original configuration. Then, because rod λ_k begins service at $y - \lambda_k$ and $y \notin B$, it must be the case that x_{n-1} lies in a busy cycle initiated by x_n . In the reduced configuration consider the assignment η that assigns rods to the rigid points $x_i, 1 \leq i \leq n - 2$, exactly as π does, and assigns rod λ_k to the free point $x_F = y - \lambda_k$. We then have rod λ_k departing at time y in the reduced configuration.

The choice of x_F and η we have created from x_n and π in this case is identical to that in sub-case (b) of the preceding case. However, recall that this case could only have occurred if x_{n-1} lies in the busy cycle initiated by x_n . Let z denote the departure time of the rod assigned to x_{n-2} when one considers *only* the rods assigned to the rigid points $x_i, 1 \leq i \leq n - 2$, according to the restriction of the assignment π (or equivalently η). The work remaining in the system at time x_{n-1} in the original configuration is then strictly less than $(z \vee x_{n-1}) - x_{n-1} + \lambda_l$. It follows that

$$y - \lambda_k - x_{n-1} < (z \vee x_{n-1}) - x_{n-1} + \lambda_l,$$

which is to say that

$$y < (z \vee x_{n-1}) + \lambda_k + \lambda_l.$$

Thus, this case cannot occur for the same values of y for which sub-case (b) of the preceding case occurs. This completes the proof of the theorem. \square

References

1. S. Asmussen, *Applied Probability and Queues* (John Wiley and Sons, 1987).
2. A. Rybko and S. Shlosman, Poisson hypothesis for information networks. I: Domain of validity, Math Arxiv, paper No. math PR/0406110.