Variance Analysis of Hall Call Response Time: Trading mean hall call response time for increased variance performance.

by Jon B. Halpern - Millar Elevator Industries Inc.

Abstract

Modern day computerized elevator dispatch systems utilizing cost function minimization techniques show that the statistical variance of the hall call response time can be decreased, making the system less random. A measure of variance performance has been derived in a prior paper called "Variance Performance Factor". Using this measure, I will investigate if there is a direct relationship that can be controlled between the mean and the variance . In other words, can mean response time be traded for increased variance performance?

Introduction

Using the analysis concepts that were developed in *Variance Analysis [1],* I will show that variance can be controlled, and that there is a correlation between the variance performance, the mean response time, and the algorithm. I chose to experiment on an existing installation rather than simulate on a computer. Obviously, a simulation is much easier to experiment on, however, the statistical process that is used to generate traffic is well known and is argued that this is not like real traffic. Real traffic creates random disturbances to the elevator system that effects statistical performance.

The elevator system used in the experiment was programmed with an algorithm to dispatch cars using a cost function that incorporated mean response time (like an ETA process) plus a variance reduction factor. I will define the cost function that was used and then show the data that was collected as we varied the controllable parameters of the cost function.

Definitions

U	=	Utilization Factor.
1/T	=	Average Hall Call Arrival Rate in Calls per Second.
R	=	Average Hall Call Response Time in Seconds.
1/R	=	Average Hall Call Response Rate in Responses per Second.
C	=	The number of cars in the system.
r_i	=	The response time to the i-th hall call.

Review of Variance Analysis

Variance Performance Factor and System Utilization Factor were introduced in an earlier paper, *Variance Analysis* [1] and discussed as a method in measuring how random a system's response is compared to a conventional dispatch system.

Within this paper, these two factors will be used extensively in the evaluation of system performance, therefore a short review on how they are calculated is necessary.

Utilization Factor

*U = Hall Call Arrival Rate / (Hall Call Response Rate * # of Cars)*

$$U = (1/T) / ((1/R) * C) = R / (T * C)$$

figure 1.

Utilization Factor provides a relative measure of how busy a system is with respect to its own average response time. As utilization approaches 1.0 the system is known as being highly utilized. Highly utilized systems should show a higher sensitivity to increases in traffic, and when it exceeds 1.0 the waiting time should increase infinitely, as in a simple queuing system.[3]

Variance Factor

V = (Variance of Response / Mean Response Squared) 100*

$$V = (Var(R) / R^2) * 100$$

figure 2.

Variance Performance Factor measures the relative performance of a system as compared to a poisson process. *Variance Analysis [1]* shows that a conventional relay elevator system behaves like a poisson process. A poisson process has a variance that is equal to its mean squared. To compare a computerized elevator system to the performance of a conventional relay elevator system, form the ratio of the new systems variance divided by its mean response squared [That is the equivalent to the conventional systems variance.]. The multiplication by 100 is just to produce a workable number where 100 is equivalent to a conventional system and factors less than 100 indicate higher performance.

Experiment

The experiment was based on an installed computerized group system in what I will call a very "regular building". I define a regular building as one that has a rectangular configuration of cars serving floors, with even floor heights, no blind shafts and all floors served by all cars with no basement. I also require that the population be fairly evenly distributed with no special services or features in the building such as cafeterias, etc.

The computerized dispatch system had a cost minimization algorithm installed that was generally simple and easy to change. The traffic gathering system was incorporated into the system and the raw data was easily collected for analysis.

Although the cost algorithm was generally simple, there are other features within the system that control several aspects of the performance. For the experiment itself the parameters of these features will be held constant.

I will show that by varying the parameters in the cost function alone, the variance factor will be affected as well as the mean response time of the system.

The Cost Function

$$J = \frac{K_1}{N} \sum_{i=1}^{N} R_i + K_2 L$$

J	=	The cost of a particular set of assignments.
N	=	The number of calls present at a given moment in time.
R_i	=	The response time of the i-th call.
K_1	=	The weighting factor of the mean response time.
L	=	The longest estimated response call in the system.
K_2	=	The weighting factor for the longest call.
N	=	The number of active hall calls registered.

figure 3.

This cost function will try to minimize the average response time as well as put additional pressure on the longest call in the system. It is a balancing act between the longest and the average with the emphasis being controlled by the weighting factors.

Although it is not readily apparent, the cost function J is dominated by the term L if both of the weighting factors are equal, because L is always greater than or equal to the mean response time. Therefore L/R_{ave} is greater than or equal to 1.0 and typically can be 4 or 5 or even as large as 10.0. This is something to keep in mind when comparing the effect of weighting factors.

Although this cost function is not optimal in minimizing the variance, it is a much easier function to compute than the actual variance . It is not my intention to design the optimal variance dispatch system for you, however I will show you the cost function for such a dispatch system.

$$J = RK_1R + K_2 \operatorname{var}(R)$$

$$\operatorname{var}(R) = \frac{1}{N-1}\sum_{i=1}^{N}(r_i - R)^2 \quad \text{where } N > 1 \qquad [2]$$

r_i = The response time to the i-th hall call.
R = The average response for a set of assignments.
$var(R)$= The variance calculated for a set R.

figure 4.

The cost function in figure 4 is much more difficult to compute and implement in a real time dynamic system because you must compute the variance for all sets of possible assignments and then pick the minimum. The number of computations needed to base a decision on this criterion is $C^{(N-1)}$, or for a 5 car group with 11 active hall calls $5^{10} = 9 \times 10^6$, a rather large number. It would take a rather powerful computer to compute that many variances and meet the requirements of a dynamic system.

Test site

The test site was a very simple 4 car group with ten stops in line . All cars serve all floors and the main entrance to the building is at the first opening. The occupancy is mixed tenancy with occasional inter-floor traffic. The floor square footage is uniform and large, 30,000 sq. ft. rentable. A typical 24 hour period can see approximately 3450 registered hall calls, most of which occur from 8:00AM till 6:00PM. In the data presented, all of the studies had 3450 +/- 250 hall calls

Test

The test consists of varying the two weighting factors K_1 and K_2 and then performing a traffic study for a given day. This process continued as we varied K_1 and K_2 for several weeks.

In order to analyze the data, I will form a ratio of K_2 / K_1 which I call the starvation ratio, where this ratio will typically be less than 1.

Results

It would be impossible for me to show any significant part of the data within this paper as there are virtually weeks of traffic analysis broken down into segments of 50, 100 and 500 calls analyzed at a time.

In figure 5 there are results for days of different starvation ratio. I have organized each set with increasing utilization rates. Each line of data contains a group of 100 calls where the utilization, average response time and variance factor were calculated. This figure is a selected set of data compiled from the complete set as an example of what was typically observed.

In figure 6 the total day averages for all performance measurements of different starvation rates are compared. This figure is the basis for all of the graphs shown, as well as a basis for the conclusions.

Analysis

The first observation is that at very low utilization factors all of the systems perform slightly worse than conventional systems with respect to the variance factor. As the starvation ratio decreases, the systems are less effected by light loads; however all of the systems have low response times under the light load condition. As the utilization increases above 20%, all the systems improve significantly over conventional systems. (See Graph 1 - Utilization vs. Variance).

The second observation is that higher starvation ratios yield higher average response times as well as lower variance factors, i.e. higher performance. This does not occur in every individual case, but when you take the average over the day this proves out rather well. (See Graph 2 - Starvation Ratio vs. Response).

Additionally, higher starvation ratios also yield higher utilization factors because the response times are increasing while the hall call arrivals are remaining constant, i.e. traffic intensity remaining constant. (See Graph 3 - Starvation Ratio vs. Utilization).

In one particular group of calls, the response time went up to 44 seconds, but the variance went down to an all time low for the study at 37. This is excellent variance performance but unacceptable mean performance.

Typical Data By Starvation Ratio

K_2/K_1	Utilization	Response	Variance Factor
1.0	12	7	141
1.0	20	9	94
1.0	36	16	52
1.0	41	15	88
1.0	54	21	77
1.0	67	21	49
1.0	90	26	43
1.0	107	44	37

K_2/K_1	Utilization	Response	Variance Factor
.667	11	6	130
.667	17	11	110
.667	30	15	67
.667	45	14	67
.667	57	21	62
.667	71	27	55
.667	85	29	53
.667	101	31	52

K_2/K_1	Utilization	Response	Variance Factor
.10	7	9	125
.10	12	8	84
.10	23	10	64
.10	21	11	92
.10	33	13	93
.10	49	14	76
.10	55	16	71
.10	74	19	70

K_2/K_1	Utilization	Response	Variance Factor
.05	8	11	105
.05	11	7	107
.05	21	11	97
.05	33	13	91
.05	36	13	66
.05	44	14	90
.05	55	16	60
.05	58	16	74

Figure 5.

Average Daily Performance for Different Starvation Ratios

K_2 / K_1	Utilization	Response	Variance Factor
3.0	64	23	66
1.0	69	24	64
.667	61	21	65
.25	55	18	71
.1	48	15	77
.05	42	14	89
.033	52	16	82
.02	48	16	81

Figure 6.

Conclusions

First, it is apparent that the variance is a controllable variable, and that by adjusting the parameters on the cost function, I can trade off response time for variance. (See Graph 4 - Response vs. Variance). The only time this is not true for the cost function is under very light traffic, i.e. low utilization rates.

Second, in graph 4 - Response vs. Variance, notice that a limit seems to have been reached for both variance performance and mean response time. The mean response time limit is controlled by the physical system constraints (speed, capacity, door times, etc.). I believe the variance performance is limited by the cost function used. Remember in the cost function that L is never less than R. If I want to maximize variance performance I would set $K_1 = 0$ then $J = K_2L$. Remember $L > R$ therefore $L=R+D$ then $J=K_2R + K_2D$ and the cost function is still dependent on R.

Therefore, what is really needed to completely control the variance is a cost function similar to the optimal system described prior, or at least the variance measure should be a simple metric measuring the distance from the actual calculated mean such as $K_2 |D|$. This algorithm would bring the system performance close to what is decribed as the ideal system, a Gaussian distribution centered on the mean response time with minimal variance[4].

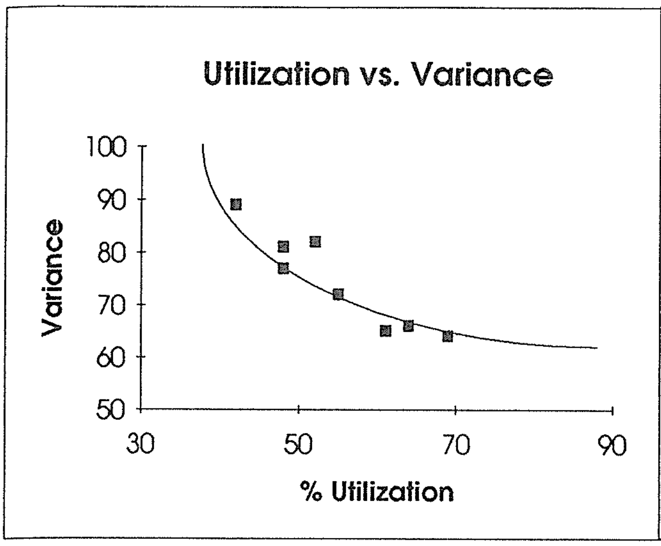

Graph 1.

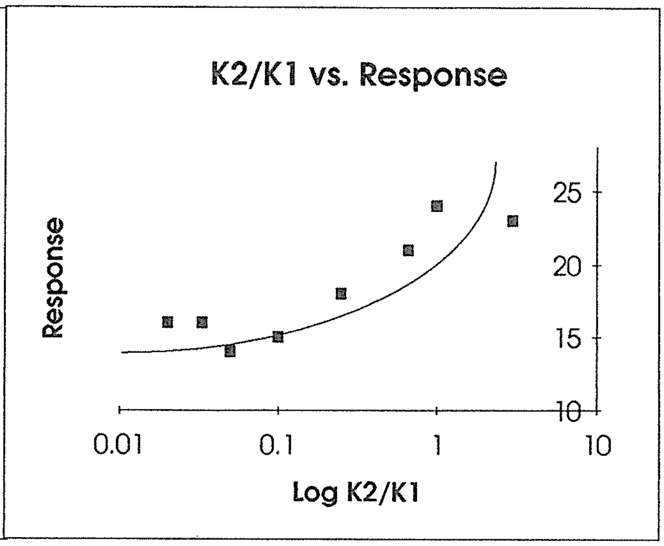

Graph 2.

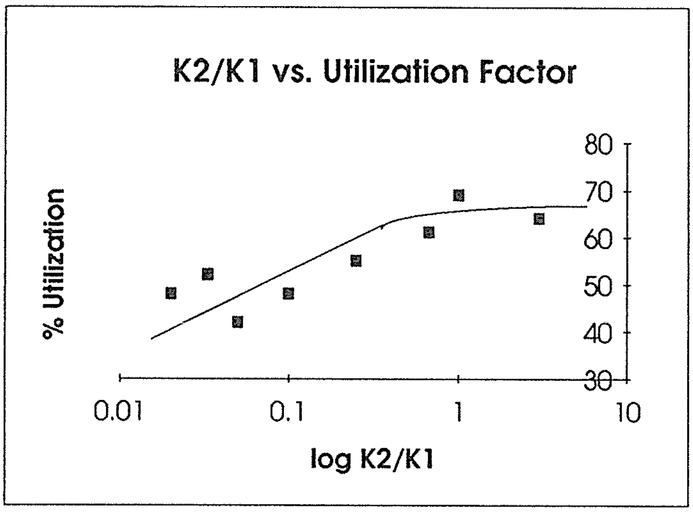

Graph 3.

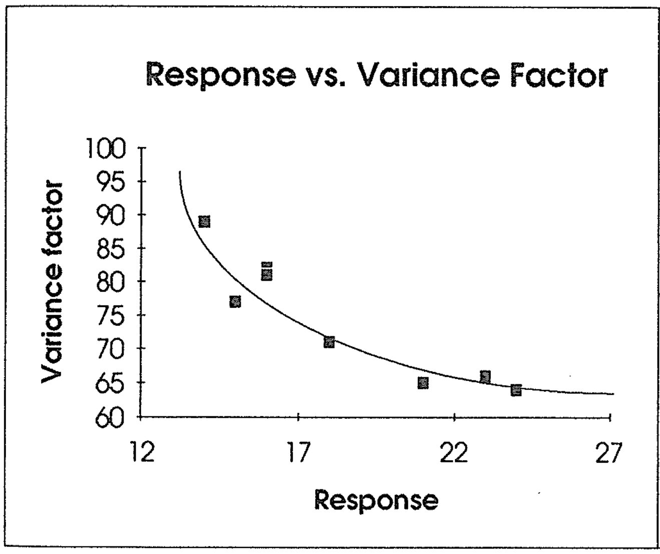

Graph 4.

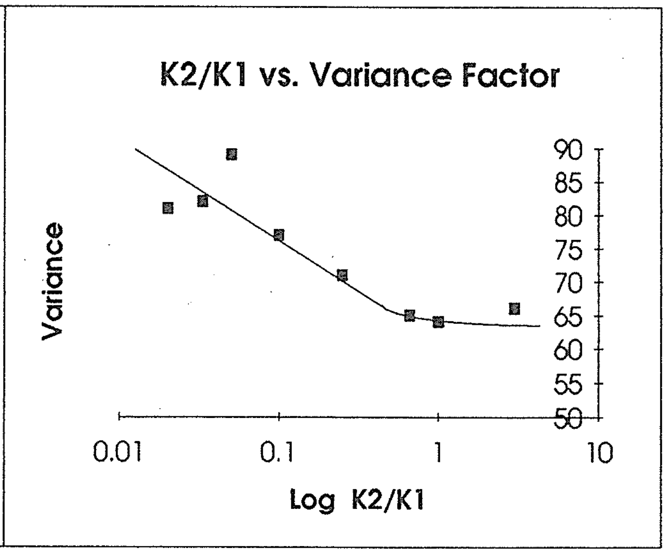

Graph 5.

Acknowledgment

I would like to thank Glenn Davis and Ricardo Abal at Millar Elevator for their help in both the experiments, modifying the sites and performing the studies.

Bibliography

[1] *Variance Analysis : A new way of evaluating elevator dispatch systems.* Elevator World, Mobile, Alabama , September 1992

[2] *Optimum Systems Control.* Andrew Sage and Chelsea White III, Prentice Hall, Englewood Cliffs, New Jersey, 1977

[3] *An Introduction to Queuing Systems 2nd Edition.* B.V. Gnedenko and I.N. Kovalenko, Birkhauser, Boston 1989.

[4] *Elevator Analysis and Design.* Revised *2nd Edition.* G.C. Barney and S.M. dos Santos, Peter Peregrinus Ltd., London, U.K. 1985

Biography

Jon Halpern is the Executive Vice President of Millar Elevator Industries Inc. in New York City. Mr. Halpern has a Bachelor of Science in Electrical Engineering from the George Washington University, 1979, a Master of Science, and a Professional Degree in Electrical Engineering from Columbia University, 1981, 1984. He has spent his entire 15 year career with Millar Elevator and is presently responsible for Modernization and Maintenance Sales, Research and Development, and Manufacturing.