

Managing the Performance of Lotus Notes: A Control Theoretic Approach

Neha Gandhi¹, Joseph L. Hellerstein², Sujay Parekh³, and Dawn M. Tilbury⁴
IBM TJ Watson Research Center²⁻³

Department of Mechanical Engineering, University of Michigan^{1,4}
{gandhin@umich.edu, hellers@us.ibm.com, sujay@us.ibm.com, tilbury@umich.edu}

The wide spread use of information technology (IT) for eCommerce and other revenue generating activities has created intense intense interest in ensuring reliable performance. Herein, we describe work done with Lotus Notes whereby an external controller is used to manipulate the SERVER_MAXUSERS tuning parameter to control the number of active users, a metric that tracks end-user perceptions of performance. Our methodology employs classical control theory and statistical models to do controller construction.

1 Introduction

Wide-spread reliance on IT systems has focused increasing attention on service level management. A commonly used approach is to take an existing **target system** and add a controller that has access to the metrics and tuning parameters of the system. Based on the feedback information present in the metrics, the controller manipulates the tuning parameters to achieve the desired service level objectives. Examples of such **closed-loop** software systems abound: the network dispatcher [6, 5], which adjusts load balancing parameters in clusters of web servers; the Multiple Virtual Storage (MVS) workload manager [1], which adjusts memory allocations and other operating system tuning parameters to achieve response time and throughput objectives; and fair share schedulers [4], which adjust Unix *nice* tuning parameter to achieve fractional allocations of CPU.

Herein our focus is Lotus Notes, a widely used workgroup and email system. Notes lacks measures for directly assessing the performance delivered to end-users. As a result, administrators often use the number of active users (those with currently executing requests) as a surrogate to gauge performance. Unfortunately, there is no tuning control that manipulates this metric directly. Rather, it must be adjusted indirectly using the SERVER_MAXUSERS parameter that controls the number of concurrent sessions (many of which are inactive at any time). This

paper demonstrates how linear control theory can be used to address the “impedance mismatch” between the active users metric and the SERVER_MAXUSERS tuning parameter.¹

Control theory has appeal since it has been used successfully in a number of engineering disciplines. The classical controller design methodology consists of two steps:

System identification: Construct a **transfer function** which relates past and present input values to past and present output values. These transfer functions constitute a model of the system.

Controller design: Based on properties of the transfer function and the desired objectives, a particular **control law** is chosen. Techniques from control theory are used to predict how the system will behave once the chosen controller is added to it.

Previous work on the application of control theoretic techniques to computer systems ([3, 7, 9, 8, 11] to name a few) has generally used first principles to

¹We focus on linear control theory because of its simplicity and wide-spread use in mechanical and electrical engineering.

perform system identification. For example, the congestion control work typically constructs state transition equations based on detailed knowledge of (or assumptions about) the protocol, workload, losses, etc. Unfortunately, there are several short-comings with a first-principles approach. First, for complex systems, it is difficult to construct a model from first principles, so often some unrealistic assumptions are made. This difficulty has been a major barrier to applying control theory to computer systems. Second, the first-principles models often employ detailed information about the target system. Since these details may change frequently (e.g., with each software release), a first-principles approach may require expert involvement on an on-going basis. This is expensive and often impractical. Third, the first-principles approach often does not address model validation. Without model validation, it is unclear how the insights obtained using control theory relate to the system being studied.

Rather than proceeding from first principles, we advocate an *empirical* approach to system identification. Here, the input and output parameters need to be identified, just as before. But rather than deriving the transfer functions based on first-principles knowledge, a linear model is constructed, and standard statistical techniques are employed to estimate the model parameters. This approach treats the system as a black box, and thus is not affected by system complexity or lack of expert knowledge. Moreover, changes in the target system can easily be accommodated by re-estimating model parameters. In this paper, we show that the approach works well for the Notes server: R^2 is no lower than 75%, and is as high as 98%.

For the controller law, we use a saturated integral controller. The behavior of such a controller is determined by one parameter, called the *gain*. Control theory tells us that the gain should be large to obtain a fast response to changing inputs, but if it is too large, it can lead to undesirable behaviors in the system, such as controller induced oscillations. The goal of the analysis, then, is to identify how large gain can be without causing these undesirable behaviors. The particular form of the models allows us to use standard techniques from control theory to perform the analysis. Our results demonstrate that there is a remarkable correspondence between the predictions made by control theory and the observed behavior of an actual Notes server. In particular, we are able to identify the feasible gain values that satisfy the control objectives.

The remainder of this paper is organized as follows. Section 2 describes the Notes server and how this target system is embedded into a closed-loop to achieve service level objectives. Section 3 details our approach to system identification. Section 4 discusses controller design and uses empirical studies to assess the accuracy of insights obtained from control theory. We provide a summary and future work discussion in Section 5.

2 Lotus Notes and Its Closed-Loop Control

This section describes relevant features of the Lotus Notes server and provides more details on how closed-loop control is obtained for this target system.

Architecturally, Lotus Notes is a client-server system. Client software converts high-level user activity (mouse clicks, etc) into remote procedure calls (RPCs) that are sent to the server. The server maintains a queue of these in-progress RPCs. Once an RPC is serviced, an entry is made in the server log, and the appropriate response is sent to the client. Clients operate in a synchronous manner – waiting for the previous request to complete before sending a new request. The client/server protocol is session-oriented. A new session is begun after a session-initiating RPC is accepted by the server. We use the term **offered load** to refer to the load imposed on the server by client requests. In the case of homogeneous clients, offered load is expressed in terms of the number of clients. Our service level metric is the length of the queue of in-progress RPC requests, hereafter just referred to as **queue length**.

The tuning parameter `SERVER_MAXUSERS` regulates the number of users allowed to access the server at any time. This is a session-level control (as opposed to packet-level RPC controls). It operates by rejecting session-initiating RPCs once the number of connected users exceeds `SERVER_MAXUSERS`. As such, this parameter has a somewhat complex effect on queue length. In particular, changing `SERVER_MAXUSERS` has no effect until a session-initiating RPC arrives, so existing sessions are not affected.

Unfortunately, we do not have direct measurements of RPC rates and queue length. Rather, we obtain these values from a measurement sensor that samples the server log at a rate of once a minute. The queue length computation is performed by counting

RPCs that were served in the previous time quantum. However, since RPCs currently waiting in the queue are not present in the log, this approach underestimates the true queue length and true RPC rates. That is, measurement is **lossy**. We can improve the approximation by delaying one or more time units before reporting the measurements since doing so allows more RPCs to complete and hence gives us a more accurate count of the RPCs that were executing.

Notes administrators are keenly interested in controlling queue length since this provides a way to manage trade-offs between response times and throughputs. Figure 1 shows how we construct a closed-loop system to control Notes queue length. Notes provides an interface for manipulating tuning parameters such as `SERVER_MAXUSERS`. We use the measurement sensor described above to obtain values of queue length. The administrator specifies a desired value, or **reference value**, for queue length. The reference value specifies a management policy that the closed loop system tries to achieve. The controller takes as input the **control error**, which is the difference between the reference value and measured value of queue length. Depending on the current (and past) values of the error, `SERVER_MAXUSERS` is adjusted. The algorithm that determines the value of `SERVER_MAXUSERS` is called the **control law**.

3 System Identification and Validation

This section describes our approach to system identification and its application to Lotus Notes. System identification has three parts. The first, block diagram construction, identifies the significant functional components and their input-output relationships. The second, transfer function formulation, models the input-output relationships of each element in the block diagram. The particular form of the models we construct (linear transfer functions) is important because it enables us to leverage a large set of analysis techniques that are available in control theory. The third, parameter estimation and evaluation, assesses the quality of the model developed.

3.1 Block Diagram Construction

A block diagram depicts the components of a system and the flow of information between them. Figure 1 provides a convenient starting point for modeling the Notes server. Here, RPC rates and `SERVER_MAXUSERS` are inputs, and queue length is the output. This is depicted in Figure 2(a).

There are two problems with the forgoing. First, it is incomplete in that lossy measurements are not considered. Thus, a separate sensor component should be included.

The second problem is more involved. Since `SERVER_MAXUSERS` has an indirect effect on RPC rates, the two inputs are not independent. Hence this is not a linear system. To address this, we divide the operating region of the Notes server into two regions. In the first, `SERVER_MAXUSERS` exceeds offered load and so the tuning parameter has no effect. In the second, `SERVER_MAXUSERS` is lower than offered load and so exactly `SERVER_MAXUSERS` users are allowed onto the system. We focus on the second operating region. Here, the offered load value is not relevant (as long as we stay in this region) and hence, it can be ignored. In other words, there is no need to consider RPC rates as an input.

Can we adequately model the Notes server if `SERVER_MAXUSERS` is the only input to the transfer function? To answer this question, Figure 3 plots queue length where the offered load is 300 users and `SERVER_MAXUSERS` increased by 20 every 20 minutes. (These data are obtained using the experimental set up described in Section 4.3.) The impact of `SERVER_MAXUSERS` is clear, suggesting that it would be sufficient. A more quantitative assessment is provided in Section 3.3.

This results in the block diagram in Figure 2(b). Note that this is a single-input, single-output system.

3.2 Transfer Function Formulation

A transfer function quantifies the relationship between inputs and outputs. Throughout, we assume that time is discrete with uniform interval sizes for a linear system with input $x(t)$ and output $y(t)$. By input-output relationships, we mean a linear model of the form

$$y(t) = \sum_{i=1}^n a_i y(t-i) + \sum_{j=0}^m b_j x(t-j) \quad (1)$$

where (n, m) is the order of the model, and the a_i, b_j are constants that are estimated from data. When values for n, m, a_i, b_j are specified, this is the transfer function of the linear system. (Technically, the transfer function is expressed as a z transform, which can be obtained in a straight-forward way from the time domain expression above.)

Now consider the Notes server. Let $q(t)$ be the actual queue length at time t and $u(t)$ be the value




Figure 1: Closed-Loop Control of Lotus Notes





Figure 2: Models of Notes server (open-loop)




Figure 3: Effect of SERVER_MAXUSERS on Queue Length

of SERVER_MAXUSERS. $m = 0$. As shown in Table 1, the following is a good fit to the data:

$$q(t) = a_1 q(t-1) + b_0 u(t).$$

This is a first order model since $\max(n, m) = 1$.

Modeling the Sensor is a bit more involved. Let $m(t)$ be the measured queue length as output by the Sensor. Once again we use a first order model: However, this is not enough since the manner in which measurements are taken introduces a delay d time units. In effect, we want $m(t)$ such that: $m(t) = m_0(t-d)$. That is,

$$m_0(t) = a_1 m_0(t-1-d) + b_0 q(t-d) + b_1 q(t-1-d).$$

3.3 Parameter Estimation and Model Evaluation

Given the forgoing functional forms, we must estimate their parameters. Our approach is statistical. First, measurements of the target system are obtained while varying the input parameters in a controlled way, such as the data in Figure 3. Then, we use least-squares regression to estimate the a_i, b_j for different values of (n, m) . In general the fit of the model improves as (n, m) are increased. We seek a model that has adequate fit and a low order.

How well do these transfer functions characterize the input-output relationships in the real system?

One way to answer this question is to use the metric R^2 , the fraction of the variability of the output variable that is explained by the transfer function model. It turns out that a first order model provides a good fit for both $q(t)$ and $m(t)$. Table 1 reports values of the a_i, b_j and R^2 for these transfer functions. For the Notes-Server, R^2 is quite large, almost 98%! This is an excellent fit. The quality of this model can be further assessed by plotting observed values of queue length versus those predicted by the model, shown in Figure 4. Note that almost all observations lie close to the line of unit slope where the predicted value equals the actual value.




Figure 4: Comparing $q(t)$ Model Predictions with Observed Values

	Delay	R^2	a_1	b_0	b_1
Notes Server	All	97.6	0.4261	0.4709	0.0
Sensor	0	75.5	0.6371	0.1692	-0.1057
	1	83.7	0.7991	0.7182	-0.6564
	2	91.2	0.9237	0.9388	-0.9092

Table 1: Model R^2 Values and Coefficients

For the Sensor transfer function, R^2 is smaller, although still acceptable. Note that as d increases so does R^2 , and a_1 approaches 1. Both effects are expected since with longer delays, measured values approach the actual value.

4 Controller Design and Assessment

Having completed system identification, the next step is to design and assess one or more controllers. We begin by describing how to construct the controller from a control law. We then use control theory to gain insights into controller behavior, especially the presence of controller induced oscillations. These predictions are assessed using measurements of a real Notes server.

4.1 Control Law and Closed-Loop Analysis

Control theory provides a systematic way to study feedback systems. Here, we show how to construct a transfer function of a closed loop system based on the transfer function of the target system in Figure 2(b). By constructing a closed loop system, we mean that the output of Figure 2(b) is fed back to the controller, which in turns compares this to the reference value $r(t)$. Based on the difference between these two values, the controller computes a new setting $u(t)$ for the control, which in our case is the value of `SERVER_MAXUSERS`.

The starting point for controller design is a control law that describes how the controller operates. We focus on **integral control** [10], a widely used technique that is a reasonable approach for the Notes server. Only one control law is considered since our objective is to demonstrate the value of our methodology.

A time-domain expression of the integral control law is

$$u(t) = u(t-1) + K_i e(t) \quad (2)$$

where $u(t)$ is the new control value at time t , and

$e(t) = r(t) - m(t)$ is the control error. The parameter $K_i > 0$ is called the *gain*. Intuitively, this control law dictates that `SERVER_MAXUSERS` be adjusted incrementally based on its previous value and the gain-weighted control error.

For an integral controller, the intuition is that higher K_i values lead to a faster response. However, care is required since larger values of K_i can cause oscillations or even instabilities.

There is a problem with directly translating this control law into software that is used for controlling Lotus Notes. Specifically, if $|K_i e(t)|$ is too large, `SERVER_MAXUSERS` is set to a value that can cause a software error. To avoid such situations, we limit the range of `SERVER_MAXUSERS` by extending the control law: $\forall t : Min \leq u(t) \leq Max$. Such **saturated controllers** are not linear. Thus, our modeling is restricted to regions in which these bounds are not reached.

4.2 Analytical Studies

This subsection uses classical control theory to evaluate the closed-loop system we study. We know from control theory that K_i should be as large as possible to provide a fast response. The issue addressed here is to predict when K_i will be so large that there are controller induced oscillations. Such a prediction is made by studying properties of the transfer function for the closed-loop system.

The approach we employ is root locus analysis. This technique examines the poles and zeroes of the frequency domain representation of the closed-loop input-output relationships. Root locus is performed by many standard software packages (e.g., MatlabTM), and so it is not discussed here. Rather, we focus on the interpretations provided by root locus.

The poles and zeros provide insight into stability and controlled-induced oscillations. If a pole of the closed loop transfer function lies outside the unit

circle, then the feedback system is unstable. That is, a bounded input produces an unbounded output. Another fact of interest relates to poles that have a non-zero imaginary part. The response due to a complex pole has a sinusoidal component that increases the variability of the output.

Root-locus plots provide a systematic way to study the location of poles in the Complex plane. Figure 5 shows root-locus plots for the closed loop system with an integral controller. Consider the left most plot, which addresses $d = 0$. The horizontal axis of the plot corresponds to real part of the pole and the vertical axis to its imaginary part. To provide a frame of reference, there is a unit circle centered at 0. The x's indicate poles, and the o's indicate zeros. The root-locus is the curve inside the unit circle that traces the poles as K_i increases from 0 to ∞ . Since all poles lie within the unit circle, there is no problem with stability. Further, observe that for small K_i (e.g. 0.1 and 1), the poles lie on the real axis. Thus, there is no sinusoidal component associated with the step response for these gains. However, for larger K_i (e.g. 5 and 50), there is a non-zero imaginary component to the poles. This suggests the presence of controller induced oscillations that increase the variance of queue length.

Now consider the root locus plot for $d = 2$, which is the right plot in Figure 5. While $K_i = 0.1$ lies on the real axis, poles for the other gains have non-zero imaginary components. Hence, we expect controller induced oscillations that result in higher variability for queue length.

We further observe that for $d = 2$, the pole at $K_i = 5$ lies outside the unit circle. This suggests a stability problem. Of course, the output of the system cannot become unbounded since we have bounded the range of values that `SERVER_MAXUSERS` is assigned. However, large gains can cause another problem—a **limit cycle** in which the tuning parameter only takes on values in $\{u_{\min}, u_{\max}\}$. We discuss this further in Section 4.3. The analysis thus reveals that if we wish to introduce a delay in order to obtain more accurate queue length information, it severely limits the range of gain values that we may use, and thus limits the responsiveness of the control system. This analysis gives us a quantitative characterization of the tradeoffs between accuracy and recency in the sensor data.

4.3 Empirical Assessments

Here, we present empirical results for various values of K_i used in an integral controller in a real system with a synthetic workload. We study how the predictions made by control theory compare with the behavior of the real system.

The testbed for our experiments consists of a workload generator, product level Notes server, a sensor running on the Notes server, and a controller running on a third machine (so as not to perturb the Notes server). The workload generator simulates the activity of multiple clients by running n copies of an identical script that sends RPCs to the server. These scripts are executed repeatedly with a one minute delay between executions. The script was selected from the NotesBench suite, a standard for such workload generation. During the experiment, the offered load to the server (i.e, the number of users trying to issue requests) is kept constant at 200 users. The reference queue length is initially set to 10, and after 60 minutes is changed to 25.

Now consider the behavior of the closed loop system if the controller is disabled. This is effected by modifying the control law so that `SERVER_MAXUSERS` is set to offered load. Figure 6 displays the result for both $d = 0$ and $d = 2$. In the former, queue length hovers around 80. In the latter, it's around 100. These results are consistent with the fact that $d = 0$ is more lossy than $d = 2$. We also see substantial variability in both cases, with changes in queue length of 40 being common. This variability comes, in large part, from having random think times between script executions.

Note further that there appears to be some oscillatory pattern in that queue length alternates between small and large values. To understand why, note that the duration of a script is also one minute. Further, recall that clients operate synchronously in that only one request can be outstanding. Thus, we tend to alternate between sets of clients that are executing and those that are waiting to execute. Thus, even without a controller modifying `SERVER_MAXUSERS`, significant variability is present.

Figure 7 shows the effect of the controller for $K_i \in \{0.1, 1, 5\}$ and delays of 0 and 2. The figure consists of 12 plots presented in two columns. The left column is $d = 0$ and the right column is $d = 2$. There are three parts to the figure. Each part considers a single value of K_i . For example, part (a) consists of the first two rows of plots. The first row

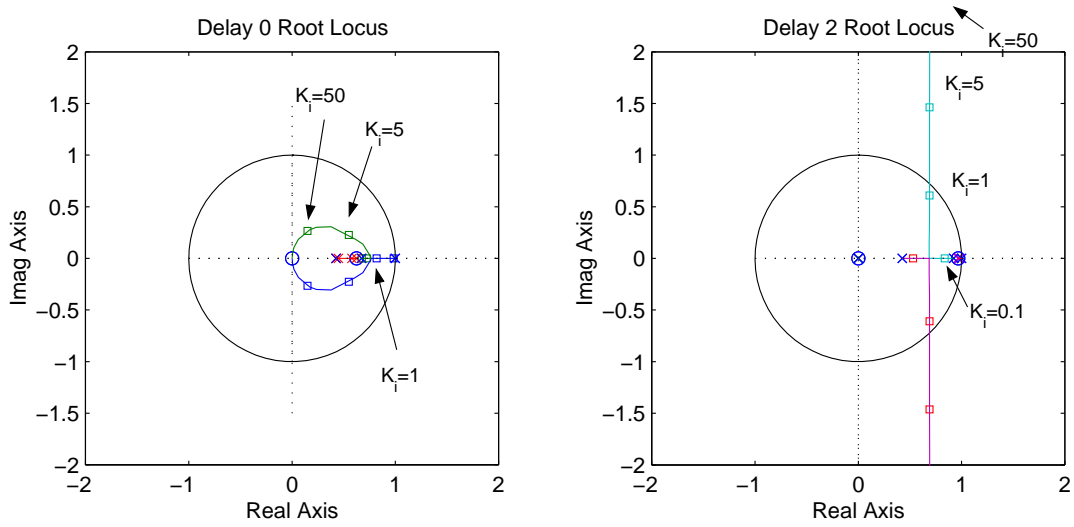


Figure 5: Root-locus plots for delay=0 and delay=2




Figure 6: Uncontrolled system run ($\text{SERVER_MAXUSERS} > \text{offered load}$)

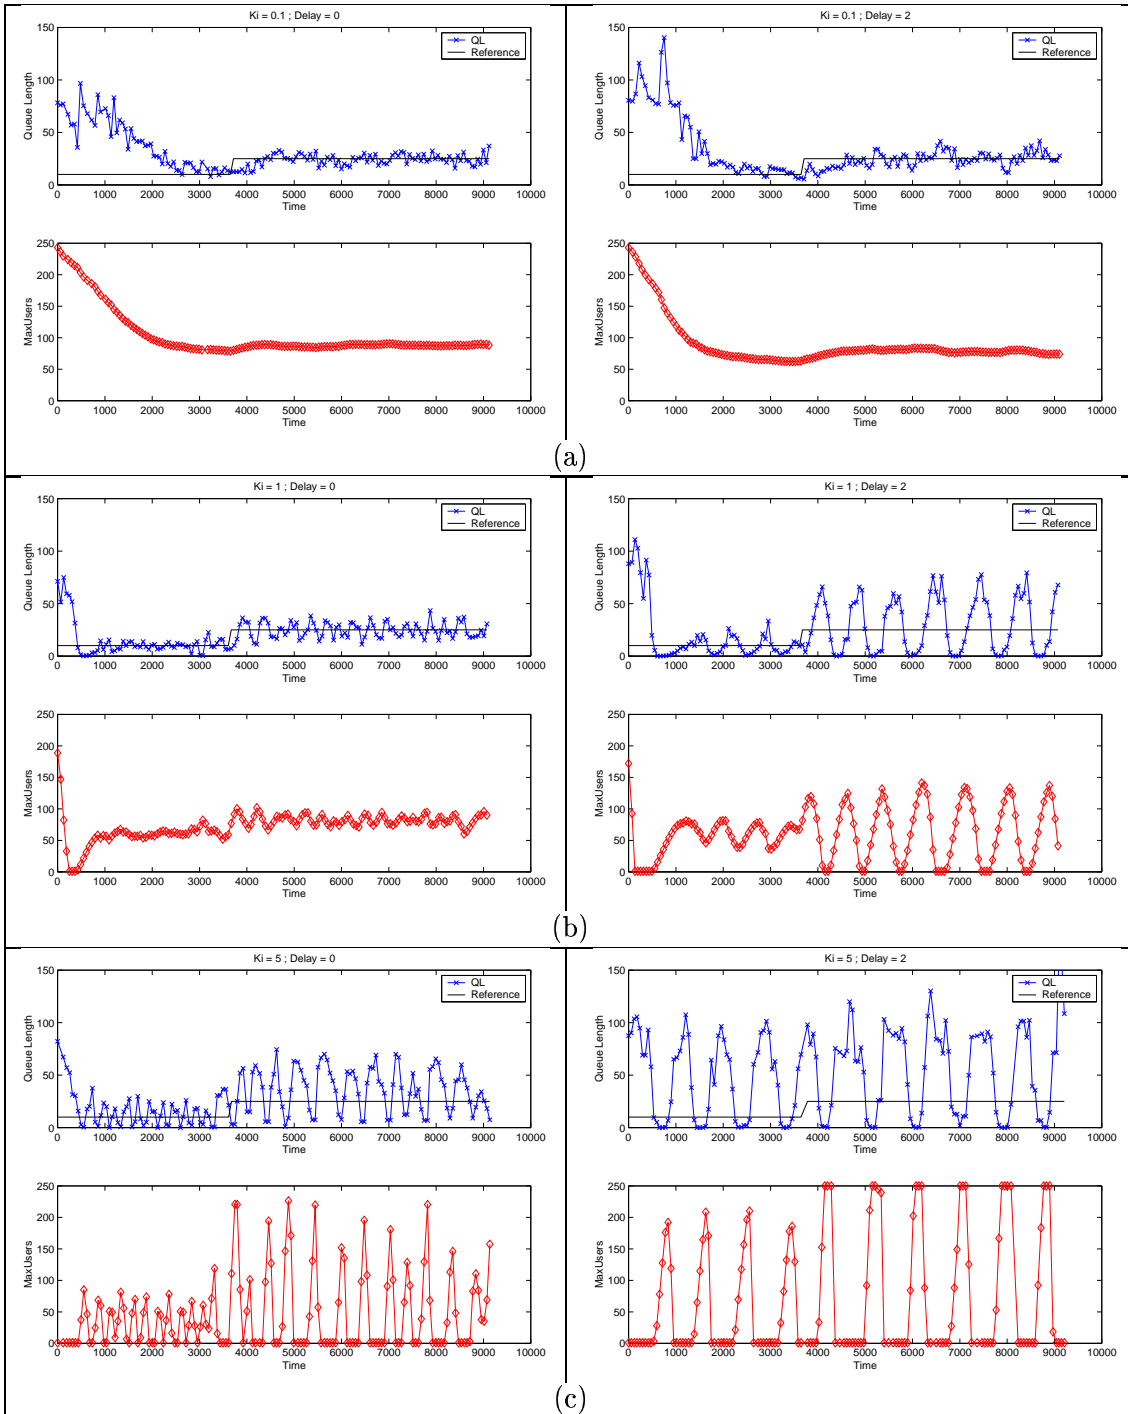


Figure 7: Effect of controller on a real system

plots queue length (and the reference value) versus time. The second row shows the value of the control at the same time as the queue length plot. Part (b) does the same for $K_i = 1$, and Part (c) displays $K_i = 5$.

Consider Part (a). While there is an initial transient, the queue length converges to the reference value. There is some variability, but variance is considerably smaller than in Figure 6. This observation holds for both $d = 0$ and $d = 2$. These results are consistent with the root-locus analysis that found the pole at $K_i = .1$ has a zero imaginary component.

Now consider Part (b) in which gain has increased by a factor of ten. We begin with $d = 0$. Here, variability is substantially larger than for $K_i = 0.1$. However, there does not appear to be a controller induced oscillation. Also, queue length values remain centered on the reference value suggesting that bias is small. The situation for $d = 2$ is not so good. There is a pronounced cycle in the queue length values, a cycle that corresponds to a cycle in the values of the control. This suggests a controller induced oscillation. We note that root-locus analysis predicted both of these results in that: (a) there is no pole with a non-zero imaginary component for $K_i = 1, d = 0$ and (b) there is such a pole for $K_i = 1, d = 2$.

What is the reason for the oscillations in queue length? Primarily, this is the result of overcompensation. That is, $K_i = 1$ is so large for the $e(t)$ in $d = 2$ that large positive errors cause the controller to increase `SERVER_MAXUSERS` and this in turn causes the next $e(t)$ to be so negative that `SERVER_MAXUSERS` greatly reduced, and so on. This is evident from the plot of $u(t)$ in Figure 7 Part (b).

In Part (c), gain is fifty times larger than in Part (a). Variability is quite large, even larger than in the uncontrolled system. Indeed, there are oscillations that are clearly related to changes in $u(t)$. Recall that root-locus analysis predicts the presence of controller induced oscillations for this case.

It is instructive to consider an extreme example, one that clearly violates the constraint of $K_i < 5$ that was established in our preliminary analysis. Figure 8 displays the results of studies done for $K_i = 50$. Here, we see a strong limit cycle for the control value, and the resultant queue length plot shows large oscillations. Indeed, changing the reference value has no apparent impact on the system's behavior. Instead, the control oscillates between extreme values

in a dysfunctional way.

Table 2 quantifies these results for the region where the reference value $r(t) = 25$. Several metrics are reported: mean queue length, standard deviation of queue length, and RMS (root-mean-square) error. Mean queue length relates to bias in that we are interested in the difference between this number and 25, the reference value. Standard deviation of queue length reflects variability in the system. Root-mean-squared (RMS) error with respect to the reference value quantifies how extreme the bias becomes.

K_i	Delay	Queue Length		RMS Error
		Mean	St. Dev.	
0.1	0	23.95	5.58	5.64
0.1	2	23.80	7.37	7.43
1	0	24.66	7.57	7.54
1	2	29.70	25.67	25.96
5	0	35.39	20.94	23.27
5	2	59.99	41.63	54.17
50	0	42.73	21.41	27.70

Table 2: Controller performance statistics

Note that standard deviations are small for those values of (K_i, d) that root-locus identified as only having real-valued poles. On the other hand, standard deviations are large for those values of (K_i, d) that do have complex poles. We also observe that for larger K_i, d , there is a problem with bias. This is indicated by large values of RMS and the difference between average queue length and the reference value of 25. This validates our analysis that for larger delays, we cannot get a fast-responding system (large K_i), otherwise the system quickly becomes unstable.

5 Summary and Future Work

Wide-spread reliance on IT services has created intense interest in automated techniques for achieving service level objectives for target systems (e.g., web servers, email servers). In order to design systems that respond dynamically based on feedback about their current state, a commonly-used approach is to create a controller that manipulates tuning parameters of the target system to achieve these objectives. However, a rigorous analysis of the behavior of such a closed-loop system is often lacking. If care is not taken in the design of the controller, then controller

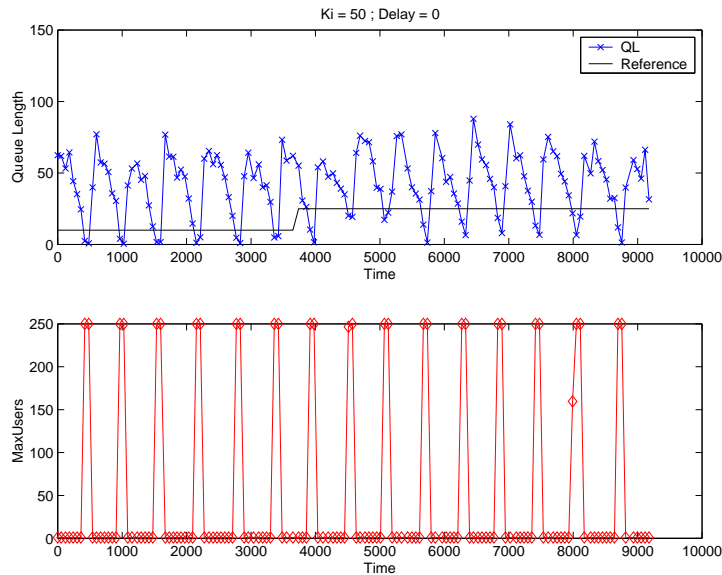


Figure 8: (Unstable) controller with high K_i

induced oscillations can arise that can degrade the quality of service delivered.

In this paper, we have demonstrated a methodology for constructing and analyzing closed-loop systems. Our starting point is classical control theory, a widely used approach in other engineering disciplines. While others have used control theory to analyze computer and communications systems, their work has not provided a general approach to modeling control systems nor has it included empirical validation. We suggest a statistical approach to system identification which is more generally applicable than the conventional first-principles approach, and also has the potential to adapt to changes in the underlying system (such as new software releases).

We have applied this methodology to the closed-loop control of a Lotus Notes server, where a controller manipulates the Notes `SERVER_MAXUSERS` tuning parameter. For system identification, we use least-squares regression to estimate the parameters of the linear model. The fit for these models is quite good: R^2 is no lower than 75%, and is as high as 98%. Interestingly, we find that it is important to model the fact that queue lengths are sampled on the server since this can introduce delays that affect controller performance.

We illustrate the value of a control-theoretic analysis by applying it to an integral controller for the Notes server, a popular technique in control the-

ory. To ensure software correctness, however, we are forced to limit the range of $u(t)$. This causes the controller to become non-linear.

With integral control, the design problem is to determine an appropriate value for the gain K_i . From control theory, we know that having a large gain makes the system more responsive. However, too large a gain can cause instabilities. Studying this using classical control theory requires that we restrict ourselves to linear regions of the controller's operation. We therefore use a simple state analysis to estimate values of the gain for which linearity should hold.

Our control-theory based analysis provides useful insights into this tradeoff about the gain value. Root-locus analysis, which is commonly used in classical control theory, allows us to predict which values of K_i cause controller induced oscillations. Our empirical studies using a real Notes system confirm these predictions. That is, in all cases where root-locus analysis predicts that a controller induced oscillation is present, this happens in our empirical studies. And in all cases where root-locus predicts that there should be no such oscillation, these oscillations are absent in our empirical results. Thus, we can choose a value for K_i that allows the system to be responsive and yet not be subject to controller-induced oscillations. Moreover, the analysis clarifies the effect of sensor-induced delays.

Note that both system identification and controller design are performed off-line, based on data that has been collected from either controlled or production runs of the target system. This allows us to use a large amount of sample data, perform more time-consuming analyses and even consult domain experts. We assume that the system evolves slowly, if at all, so the model does not need to be estimated often. An online changepoint detection [2] scheme can be employed to actively monitor the system and trigger the parameter re-estimation when required. Online adaptation of the target that is within the bounds of the estimated model is performed by the controller.

Much work remains. Our approach to identifying linear regions of operation is approximate at best. A better approach would employ describing functions, a technique used in non-linear control theory. In this paper we have restricted ourselves to a simple control law in order to demonstrate the value of this approach. We plan to study more complex controllers to assess if control theory provides useful insights as to their operation. More broadly, we are interested in applying our methodology to other service level management situations both to refine our methodology and to assess its value.

References

- [1] J. Aman, C. K. Eilert, D. Emmes, P. Yocom, and D. Dillenberger. Adaptive algorithms for managing a distributed data processing workload. *IBM Systems Journal*, 36(2), 1997.
- [2] M. Basseville and I. Nikiforov. *Detection of Abrupt Changes: Theory and Applications*. Prentice Hall, 1993.
- [3] Dah-Ming Chiu and Raj Jain. Analysis of the increase and decrease algorithms for congestion avoidance in computer networks. *Computer Networks and ISDN systems*, 17(1), June 1989.
- [4] Raymond B. Essick. An event-based fair share scheduler. In *Proceedings of ACM USENIX*, pages 147–161, Winter 1990.
- [5] G. Hunt, G. Goldszmidt, R. King, and R. Mukherjee. Network dispatcher: A connection router for scalable internet services. In *Proceedings of the 7th International World Wide Web Conference*, April 1998.
- [6] Arun Iyengar, Jim Challenger, Daniel Dias, and Paul Dantzig. High-performance web site design techniques. *IEEE Internet Computing*, 4(2):17–26, February 2000.
- [7] Srinivasan Keshav. A control-theoretic approach to flow control. In *Proceedings of ACM SIGCOMM '91*, September 1991.
- [8] Baochun Li and Klara Nahrstedt. Control-based middleware framework for quality of service applications. *IEEE Journal on Selected Areas in Communication*, 1999.
- [9] Saverio Mascolo. Classical control theory for congestion avoidance in high-speed internet. In *Proceedings of the 38th Conference on Decision & Control*, December 1999.
- [10] Katsuhiko Ogata. *Modern Control Engineering*. Prentice Hall, 3rd edition, 1997.
- [11] David C. Steere, Ashvin Goel, Joshua Gruenberg, Dylan McNamee, Calton Pu, and Jonathan Walpole. A feedback-driven proportion allocator for real-rate scheduling. In *Proceedings of Operating Systems Design and Implementation (OSDI)*, February 1999.