

Statistical Effects of Control Parameters on Throughput of Window-Based Transport Method

Qishi Wu, Nageswara S.V. Rao
Computer Science and Mathematics Division
Oak Ridge National Laboratory
Oak Ridge, TN 37831-6364
Email: {wuqn,raons}@ornl.gov

S. Sitharama Iyengar
Department of Computer Science
Louisiana State University
Baton Rouge, LA 70803
Email: iyengar@bit.csc.lsu.edu

Abstract—In window-based transport methods for stabilizing and/or maximizing the goodput at the destination, it is very important to understand the statistical properties of the transport control and performance response parameters. Based on traffic measurements collected over the Internet during a 6-month period, we formulate and test hypotheses on the main effects of two control parameters on the goodput response and the interaction effects between them. We infer from the statistical analysis that the congestion window and sleep time parameters strongly interact with each other, and they both have significant main effects on the destination goodput. Consequently the underlying randomness in network traffic must be explicitly accounted for in the design of flow control methods.

I. INTRODUCTION

Next generation network applications require transport capabilities that have not been traditionally addressed by the current implementations or analytical methods. For example, real-time remote control over wide-area networks and coordinated remote visualization of distributed datasets require “control channels” implemented over wide-area networks. In practice, these channels require a steady flow rate typically at a level much below the peak link rate. While the Transmission Control Protocol (TCP) performs very well in supporting several Internet transport tasks, it lacks the stable dynamics needed in the above applications. Thus there has been a renewed interest in developing newer classes of transport protocols. However, many of the analytical results in transport protocols rely on deterministic control methods that typically assume smooth utility functions [1]. One of the open questions is whether or not the statistics of connections be explicitly accounted for in the protocol design. This is an important consideration since the deterministic flow controls that often rely on a (suitably chosen) fixed step size do not stabilize in presence of randomness in the network delays. In fact, the step sizes have to be dynamically adapted as per the Robbins-Monro (RM) condition to achieve stable goodput over wide-area connections in recent stabilization protocols (see [2]).

The flow control mechanism in TCP employs a sliding window technique. Recent protocols such as Tsunami and SABUL [3], [4] employ an alternative rate control in which the sender restricts the congestion window to one packet and adjusts the inter-packet delay to spread the packets out evenly. We consider generic window-based control methods

of which the above two mechanisms are special cases. There are two adjustable parameters, namely the *congestion window* and *sleep time*, that control the source transmission rate, to which the destination goodput (rate of packets received without counting the duplicates) and loss rate respond. But the underlying response function is complicated because of the randomness in network delays and packet losses due to the competing traffic on the shared network links. To account for the underlying randomness in the flow control design, it is necessary to understand the statistics of the two control parameters as well as the performance responses. In this paper, we apply rigorous statistical methods to investigate how these two control parameters affect the transport performance. While the window-based protocols have been studied for decades, our analysis provides valuable insights into the statistical aspects of their parameters, and in particular indicates that simple deterministic methods are unlikely to meet the requirements.

We investigate the *main effects* which capture the differences in the mean goodput response as we change the values of either congestion window or sleep time individually. We also consider the *interaction effects* that capture the differences in the main effect goodput responses for one factor as we vary the other across the range. We formulate a generalized linear statistical model relating the control parameters to the destination goodput, and test the hypotheses based on the measurements collected over a long time span. Our statistical analysis results establish the following in the Internet environments with regular traffic:

- 1) Congestion window and sleep time strongly interact with each other in deciding the network transport performance.
- 2) Within the range of each control variable both congestion window and sleep time have significant main effects on the destination goodput.
- 3) Within the variable ranges for the test Internet link, the errors of the “noise-corrupted” measurements of the destination goodput response confined within each cell determined by a unique combination of congestion window and sleep time are identically and normally distributed with a zero mean and the same variance.

In summary, these results indicate that in general the statistical

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE OCT 2003	2. REPORT TYPE	3. DATES COVERED 00-00-2003 to 00-00-2003			
4. TITLE AND SUBTITLE Statistical Effects of Control Parameters on Throughput of Window-Based Transport Method		5a. CONTRACT NUMBER			
		5b. GRANT NUMBER			
		5c. PROGRAM ELEMENT NUMBER			
6. AUTHOR(S)		5d. PROJECT NUMBER			
		5e. TASK NUMBER			
		5f. WORK UNIT NUMBER			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Louisiana State University, Department of Computer Science, Baton Rouge, LA, 70803		8. PERFORMING ORGANIZATION REPORT NUMBER			
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)			
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)			
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES 12th International Conference on Computer Communications and Networks (ICCCN'03), Dallas, Texas, October 2003					
14. ABSTRACT see report					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified	Same as Report (SAR)	4	

effects of network delays and packet losses are significant in the window-based flow control methods, and must be accounted for in the design. Such approaches have been used for goodput stabilization and maximization in [2] based on dynamic and simultaneous perturbation stochastic approximation methods, respectively.

II. WINDOW-BASED NETWORK TRANSPORT CONTROL

We consider a window-based network transport control method that uses both User Datagram Protocol (UDP) and TCP. The sender or source generates the data and sends it to the receiver (or destination) as UDP datagrams. The sender explicitly informs the receiver of the initialization and termination of the data transmission process via TCP. This simple window-based flow control scheme has two parameters, the *congestion window* and *sleep time*, both of which control the source sending rate, typically with different effects on the throughput. The *congestion window* $W_c(t)$ represents the number of UDP datagrams that can be sent continuously as fast as the computer and communication hardware resources (CPU, memory, NIC speed, channel bandwidth, etc.) allow before waiting for a certain time period. The *sleep time* or *idle time*, denoted by $T_s(t)$, represents the amount of time the sender suspends right after sending a full congestion window of UDP datagrams in a burst until the next burst.

Upon termination, the receiver at the destination node turns off timer and calculates average goodput $g_D(t)$ as well as loss rate $l_D(t)$. Based on the above flow control model, the instantaneous source rate $r_S(t)$ is:

$$r_S(t) = \frac{W_c(t)}{T_s(t) + T_c(t)} = \frac{1.0}{\frac{T_s(t)}{W_c(t)} + \frac{1.0}{BW}} \quad (1)$$

where $T_c(t) = \frac{W_c(t)}{BW}$ is the time spent on continuously sending a full congestion window of UDP datagrams, which is determined by the congestion window size and communication hardware resources, and the system bandwidth BW , i.e. the maximum speed at which the sender can generate the signal and put it on wire. According to Eq (1), we may control the source rate $r_S(t)$ by adjusting either congestion window $W_c(t)$ or sleep time $T_s(t)$ individually, or both simultaneously. Solutions to the transport control problems for different throughput requirements may involve a dynamic adaptation of $W_c(t)$ and $T_s(t)$ to achieve the optimal source rates.

III. STATISTICS OF CONTROL PARAMETERS

Since the network traffic measurements are inherently statistical in nature, we now formulate an appropriate statistical model for studying the two control parameters. In the window-based control model, the two factors, congestion window and sleep time, are simultaneously applied in controlling the datagram transmission at the source node, which constitutes a *two-factor experiment*, as is known in statistics literature. The *main effects* correspond to the differences in the mean goodput response across the various levels of either parameter viewed individually. The *interaction effects* correspond to the

differences of main effect goodput responses for one factor as the values of the other are varied.

We apply a *general linear model* (GLM) to the random-effects model of two factors, namely congestion window and sleep time. GLM assumes that any particular observation value can be accounted for by summing up a number of predictor components and a residual term. The linearity of GLM does not manifest a linear relationship between the response and condition variables, but establishes the additivity of these components, as defined by:

$$g_{ijk} = \mu + c_i + s_j + (cs)_{ij} + \varepsilon_{k(ij)} \quad (2)$$

where

$g_{ijk}, k = 1, 2, \dots, n$: k -th observed value of the destination goodput $g_D(t)$ under the combinatorial treatment defined by the i -th level of congestion window and j -th level of sleep time.

μ : reference value, calculated as the sum of all observation values divided by the total number of observations.

$c_i, i = 1, 2, \dots, a$: main effects of congestion window on goodput response, calculated as the difference between the mean response of the subpopulation comprising the i -th level of congestion window and the grand mean μ .

$s_j, j = 1, 2, \dots, b$: main effects of sleep time on goodput response, calculated as the difference between the mean response of the subpopulation comprising the j -th level of sleep time and the grand mean μ .

$(cs)_{ij}$: interaction effects between congestion window and sleep time, calculated as the difference between the mean goodput response in the subpopulation defined by the combination of the factor levels of c_i and s_j , and the mean goodput response when there only exist main effects of either c_i or s_j .

$\varepsilon_{k(ij)}$: random error. The parentheses around subscript variable i and j can be considered as a cell defined by i -th level of congestion window and j -th level of sleep time. The random error $\varepsilon_{k(ij)}$ represents the variation among n goodput observations nested in the cell (i, j) , and is calculated as the difference between k -th observation g_{ijk} nested in cell (i, j) and the cell mean $\bar{g}_{(ij)}$.

We have the following assumptions for variables in Eq (2).

- 1) The grand mean μ is a constant.
- 2) All random variable $c_i, s_j, (cs)_{ij}$, and $\varepsilon_{k(ij)}$ are statistically independent.
- 3) Main effects $c_i, i = 1, 2, \dots, a$ are identical and independent variables, with a zero-mean normal distribution, i. e. iid $\sim N(0, \sigma_c^2)$.
- 4) Main effects $s_j, j = 1, 2, \dots, b$ are identical and independent variables, with a zero-mean normal distribution, i. e. iid $\sim N(0, \sigma_s^2)$.
- 5) Interaction effects $(cs)_{ij}$ are identical and independent variables, with a zero-mean normal distribution, or equivalently, $(cs)_{ij}$ iid $\sim N(0, \sigma_{cs}^2)$.
- 6) Random errors $\varepsilon_{k(ij)}$ have a mean of zero and a common variance equal to σ_ε , or equivalently, $\varepsilon_{k(ij)} \sim N(0, \sigma_\varepsilon^2)$.

The first five assumptions made above on the grand mean and random variables are the standard procedure required by the

GLM. However, the assumption on the error distribution is an esoteric point and must be tested to verify the analysis. This test is performed later together with other hypotheses on main and interaction effects. Taking variance of Equation (2) results in the following:

$$\begin{aligned} \text{Var}(g_{ijk}) &= \text{Var}[\mu + c_i + s_j + (cs)_{ij} + \varepsilon_{k(ij)}] \\ &= \sigma_c^2 + \sigma_s^2 + \sigma_{cs}^2 + \sigma_\varepsilon^2 \end{aligned} \quad (3)$$

From Equation (3) we know that there are four variance components in the model we constructed, which can be estimated using the mean square of error, mean square of main effects as well as mean square of the interaction effects.

IV. EXPERIMENTAL SETUP

The window-based flow control mechanism is implemented at Louisiana State University (LSU) and at Oak Ridge National Laboratory (ORNL). At the time of experiment, ORNL is connected to ESnet, which peers with Abilene network in New York. Abilene runs from New York via Washington DC and Atlanta to Houston, where it connects to LSU via a regional network. In terms of network distance, these two sites are separated by more than two thousand miles, and both ESnet and Abilene backbones as well as the hosts have significant network traffic.

The client at ORNL generates a message of a certain size and sends it to the server at LSU as a set of UDP datagrams at a fixed rate at a time for multiple times with different sending rates. We maintain a constant sending rate of UDP datagrams during each run of the message transmission by fixing both congestion window and sleep time in the flow control mechanism.

At the source node, we compute the average source rate $r_S(t)$ as the number of sent UDP datagrams divided by the transmission duration. At the destination node, the average goodput $g_D(t)$ and loss rate $l_D(t)$ are measured as the number of successfully arriving UDP datagrams and the number of lost UDP datagrams respectively, both of which are divided by the time elapsed since the first datagram is delivered.

The various source transmission rates are achieved through the independent adjustments made on either congestion window or sleep time. Particularly, we conduct the transport control experiment by varying the congestion window from 1 to 100 at a step of 5 UDP datagrams and by varying the sleep time from 1 to 100 at a step of 5 milliseconds independently. Consequently, the congestion window and sleep time have the same number of levels $a = b = 100/5 = 20$. For each particular combination of congestion window and sleep time, $n = 5$ observations are made. All the network traffic measurements are collected between the client at ORNL and the server at LSU. We repeat the whole set of experiment, i.e. multiple runs of the message transmission at various fixed sending rates, over hours, days, weeks, and months to extensively investigate the influence of control parameters on the network performances.

From the measurements collected over 6 months, we observed that the network traffic exhibits a very similar two-phase pattern. In the first phase, there is a trend of monotonic increase in goodput $g_D(t)$ as sending rate $r_S(t)$ is increased while the loss rate $l_D(t)$ remains at an extreme low level. After the sending rate reaches a certain transition point, the system enters the second phase where the goodput $g_D(t)$ starts suffering irregular decrease due to congestion collapse indicated by the high datagram loss rate $l_D(t)$. This overall behavior is quite stable although the transition position dividing these two phases and the goodput shape may slightly vary over time in the presence of diverse background traffic. We also observed that the plots of goodput and loss rate are quite non-smooth because of the dynamic network conditions that induce randomness into packet delays and losses.

This type of overall goodput response is well known but often only smoother versions are considered which correspond to a deterministic formulation [1]. However, in practical applications, when source rate $r_S(t)$ is fixed at r , the goodput $g_D(t)$ is a random variable, which is jointly distributed with distribution $G(g_D(t), r)$. By running the experiment with the same control parameters for a sufficient large number of times and computing the mean values of the destination goodput $g_D(t)$, we obtain the expected value of the destination goodput corresponding to the fixed sending rate r :

$$M_D(r) = E[g_D(t)|r_S(t) = r] = \int g_D G(g_D, r)$$

where $G(\cdot, r)$ is an unknown distribution function of real-valued random variable g_D at a given constant sending rate r . We refer to $M_D(r)$ as the *destination goodput response regression*. The long-time-span Internet measurements show that $M_D(r)$ experiences very slight variation in presence of on-host or off-host background traffic during most days.

V. STATISTICAL ANALYSIS AND APPLICATIONS

A. Hypothesis Test

In order to explore the network performance pattern and the statistical nature of the network traffic, a series of Internet experiments have been conducted over a time span of 6 months. The collected measurements are used as the data sets for the experimental statistics methods described above. We use SAS, the most common large-scale data analysis software package, to perform the ANOVA (Analysis of Variance) analysis of the random-effects two-way factorial experiment. We have the following hypotheses of interests and the test of each hypothesis is given accordingly.

- 1) *Hypothesis of main effects of congestion window on goodput response:* $H_0 : \sigma_c^2 = 0; H_1 : \sigma_c^2 \neq 0$. In random effects model, the mean square of the interaction effects is used as the error term for statistic test of main effects. Therefore, the F value is computed as: $F^* = \frac{MS(\text{congestionwindows})}{MS(\text{interaction})} = 17.91$ with numerator degree of freedom $ndf = 19$ and denominator degree of freedom $ddf = 361$. The corresponding p -value is less than 0.0001, based on which we reject H_0 . In

other words, the observation data strongly indicates that the congestion window has a significant effect on the destination goodput.

- 2) *Hypothesis of main effects of sleep time on goodput response:* $H_0 : \sigma_s^2 = 0; H_1 : \sigma_s^2 \neq 0$. Same as above, the mean square of the interaction effects is used as the error term for statistic test of main effects, and the F value is computed as: $F^* = \frac{MS(sleep\ time)}{MS(interaction)} = 5.26$ with numerator degree of freedom $ndf = 19$ and denominator degree of freedom $ddf = 361$. The corresponding p -value is less than 0.0001, based on which we reject H_0 . That is to say, the observation data strongly indicates that the sleep time also has a significant effect on the destination goodput.
- 3) *Hypothesis of interaction effects between congestion window and sleep time on goodput response:* $H_0 : \sigma_{cs}^2 = 0; H_1 : \sigma_{cs}^2 \neq 0$. The mean square of error MSE is used as the error term for statistic test of interaction effects. Therefore, the F value is computed as: $F^* = \frac{MS(interaction)}{MS(error)} = 10.35$ with numerator degree of freedom $ndf = 361$ and denominator degree of freedom $ddf = 1600$. The corresponding p -value is less than 0.0001, based on which we reject H_0 . In other words, the observation data strongly indicates that there is a significant interaction between the congestion window and sleep time.
- 4) *Hypothesis of residual normality under linear statistical model:* We use Shapiro-Wilk method to test the residual normality. The SAS program calculates the statistic $W = 0.875659$, and the corresponding p -value ($Pr < W$) is less than 0.0001, based on which, we infer that the residual normality assumption is valid in the light of measurement data.

B. Design of New Protocols

The above analysis provides insights into the significance of control parameters on the network transport performance and reveals the stochastic nature of network traffic over wide-area networks. We construct stochastic approximation models based on this statistical analysis to design new transport protocols for goodput stabilization [5] and maximization [2].

In the goodput stabilization problem, the transport control objective is to dynamically adjust the source rate such that the destination goodput is stabilized at a desired level, which is usually much lower than the maximum achievable goodput. In [5], we design a class of transport control protocols based on dynamic Robins-Monro Stochastic Approximation method for goodput stabilization. Either congestion window or sleep time are used in the protocol design and the equivalent main effects of both control parameters on the goodput response are justified by the experiment results.

In the goodput maximization problem, we aim to dynamically control the source rate to achieve high bandwidth utilization by maximizing the individual throughputs from an overall perspective. The goodput maximization problem identifies two considerations: fair share and high utilization of

bandwidth. Fair share requires that all concurrent data streams sharing the same link be equally treated, and high utilization requires that the link bandwidth is exploited to the greatest possible advantage. In [2], we design a class of transport control protocols based on dynamic Simultaneous Perturbation Stochastic Approximation method, which controls both congestion window and sleep time simultaneously, to achieve maximum individual goodput.

Since the results of previous section indicate significant randomness in the underlying process, it is insufficient to use deterministic approaches; the latter can be shown to stabilize under suitably chosen but fixed step sizes. But such step sizes cannot be shown to stabilize in presence of randomness in the control loop, and in fact did not stabilize in our implementations [2]. In the above two flow controllers, the step sizes of control parameters are adapted according to the conditions of the stochastic approximation methods. Intuitively, the step sizes have to approach zero eventually but only at a rate controlled by certain upper and lower bound conditions (see [2] for details). The experiment results show that these new transport protocols have superior performances than the default TCP in stabilizing or maximizing goodput over wide-area networks.

VI. CONCLUSIONS

Our results indicate that the statistical effects must be explicitly taken into account in the design of the window-based transport protocols, and the traditional methods based on smooth deterministic control methods might not be adequate to achieve goodput stabilization and maximization. There are a number of future directions to pursue. It would be interesting to see if the statistical results of this paper will be similar under a non-linear model. It would also be of interest to investigate flow control methods that explicitly utilize the results of this analysis in adapting the control parameters. Also, the statistical analysis of rate-based flow control methods [6] will be of interest if they can be designed within the framework of stochastic approximation.

ACKNOWLEDGMENT

This research is sponsored by High-Performance Networking Program of the Office of Science, U.S. Department of Energy, under Contract No. DE-AC05-00OR22725 with UT-Battelle, LLC, Network Modeling and Simulation Program of Defense Advanced Research Projects Agency, and National Science Foundation.

REFERENCES

- [1] S. Low, F. Paganini, and J. Doyle, "Internet congestion control," *IEEE Control Systems Magazine*, February 2002.
- [2] Q. Wu, "Control of transport dynamics in overlay networks," Ph.D. dissertation, Louisiana State University, 2003.
- [3] "Tsunami," <http://newsinfo.iu.edu/news/page/normal/588.html>.
- [4] "Sabul: A high performance data transfer protocol," <http://www.rgrossman.com/pdf/sabul-hdtp-11-02.pdf>.
- [5] Q. Wu and N. S. Rao, "Network transport daemons for stable flow channels over wide-area networks," in *Proceedings of Hawaii International Conference on System Sciences*, January 2004.
- [6] L. Peterson and B. Davie, *Computer Networks: A Systems Approach*. Morgan Kaufman, 2000.