A Dynamic Regulation and Scheduling Scheme for Real-Time Traffic Management

Steve Iatrou and Ioannis Stavrakakis, *Senior Member, IEEE*

***Abstract*—Typical rate-based traffic management schemes for real-time applications attempt to allocate resources by controlling the packet delivery to the resource arbitrator (scheduler). This control is typically based only on the characteristics of the particular (tagged) traffic stream and would fail to optimally adjust to non-nominal network conditions such as overload. In this paper, a dynamic regulation and scheduling (dynamic-R&S) scheme is proposed whose regulation function is modulated by both the tagged stream's characteristics and information capturing the state of the coexisting applications as provided by the scheduler. The performance of the proposed scheme—versus an equivalent static one—is investigated under both underload and overload traffic conditions. The substantially better throughput/jitter characteristics of the dynamic-R&S scheme are established.**

***Index Terms*—Delay variance, dynamic policy, QoS, regulation, scheduling, throughput.**

I. Introduction

A secure solution to the problem of guaranteeing the QoS of real-time applications will typically require the reservation of the maximum amount of needed resources. Real-time traffic is bursty in its nature, and therefore guaranteeing QoS leads to severe network underutilization. Alternative solutions are being considered based on "overallocation" of resources to a group of applications (multiplexing). Grouping of applications and "overallocation" of resources are the key aspects of nondegenerate statistical multiplexing.

Because of the stringent QoS requirements of real-time applications, it is expected that "traditional" statistical multiplexing schemes, such as FCFS, will not be effective for such applications. It is well understood that some tighter control should be exercised on input and output as well as in the internal processes of a multiplexing scheme which impact on its efficiency. Sophisticated call admission schemes or other types of "weak" resource reservation can provide control at the larger time scale. Proper traffic regulation and service scheduling mechanisms can provide control at the smaller time scale. Using such control schemes, statistical multiplexing will potentially provide for increased network utilization while delivering the more stringent QoS. This paper is focused on the control approaches at the smaller time scale, namely, regulation and scheduling.

In principle, a target value of a QoS metric (cell loss/dealy) may be possible to achieve through either tight traffic regulation (rate-based approach) or sophisticated scheduling (scheduler-based approach) only. In most practical cases though, some scheduling will be needed to resolve transmission conflicts among rate-based controlled applications. Similarly, some traffic filtering (regulation) will be needed to eliminate extreme traffic realization which would be hard to manage even by a sophisticated scheduler under a scheduler-based approach.

Substantial effort has been directed toward the development of regulation and scheduling schemes for real-time applications. Examples of regulation and scheduling schemes for real-time applications include delay-earliest due date (D-EDD) [1]; jitter earliest due date (J-EDD) [2], [3]; hierarchical round robin (HRR) [4]; stop and go queuing (S&G) [5]; weighted fair queuing (WFQ) [6]; packet generalized processor sharing (PGPS) [7]; rate controlled static priority (RCSP) [8]; leave in time (LIT) [9]; multirate traffic shaping (MRTS) [10]; and virtual clock (VC) [11].

Schemes based on the round robin (RR) idea, such as WFQ, HRR, PGPS, and VC, are mainly concerned with traffic isolation. Such schemes employ scheduling decisions based on traffic rates rather than packet by packet metrics. Similar metrics are employed by S&G and MRTS. The proposed policy's metrics are more comparable with D-EDD, J-EDD, or RCSP. Specifically, the traffic-management scheme for real-time applications investigated in this work may be viewed as an enhancement of the RCSP scheme proposed in [8].

An apparent drawback of schemes such as the RCSP is that the regulator and scheduling functions are separated. It is expected though that the throughput/jitter of a regulated tagged stream will be substantially modulated at the scheduler by the cumulative activity of the coexisting traffic streams. As a consequence, the effectiveness of the regulation and scheduling (R&S) scheme may be compromised significantly. This problem can be addressed to some extent by *dynamically* adjusting the regulator behavior based on state information fed back from the scheduler to the regulator. Such a dynamic policy is investigated in this paper.

In Section II, the proposed dynamic R&S (dynamic-R&S) policy is motivated and presented, along with the equivalent static R&S (static-R&S) policy, on which a comparative study will be based. In Section III, the behavior of the tagged regulator is investigated under both policies. In Section IV, the behavior of the scheduler is studied under underload conditions at the

Manuscript received April 11, 1997; revised December 1, 1998; approved by IEEE/ACM Transactions on Networking Editor S. K. Tripathi. This work was supported in part by the Defense Advanced Research Project Agency under Grant F49620-93-1-0564, monitored by the Air Force Office of Scientific Research, and in part by the National Science Foundation under Grant NCR-9628116.

S. Iatrou is with Andersen Consulting, Sophia Antipolis 06902 France (e-mail: steve.iatrou@ac.com).

I. Stavrakakis is with the Department of Informatics; University of Athens, Athens 15784 Greece(e-mail: istavrak@di.uoa.gr).

Publisher Item Identifier S 1063-6692(00)01439-4.

scheduler. In Section V, the throughput/jitter performance of the policies is investigated under overload conditions at the scheduler. The analysis and simulation study results are presented in Sections VI and VII, respectively.

II. The Proposed Dynamic R&S Policy

The typical primary objective in regulating real-time traffic stream within the network is to control jitter or the instantaneous rate (throughput). This is achieved in the RCSP mechanism [8] by enforcing a minimum spacing at the output of the regulator associated with the traffic stream of interest (tagged traffic stream). Under the RCSP scheme, each traffic stream passes through a regulator, which restores the traffic, completely or partially, based on the traffic description and the type of regulator used. The restored traffic is handed over to the respective priority queue and is scheduled in FCFS order. The rate jitter regulator employed in [8] is based on cell eligibility times (ET) defined as follows: $ET_1 = AT_1$; $ET_k = \max\{ET_{k-1} + T + \tau_k, AT_k\}$, $k > 1$, where AT denotes the cell arrival time and τ_k is a term used to provide the average rate; subscripts indicate packets. Rate jitter is controlled with respect to the ET of the previous packet of the same connection. T is the minimum cell inter arrival time specified by the source. The idea is to hold cells so that minimum inter departure time be enforced. Fig. 1 (without the feedback loop) shows a block diagram of an architecture implementing a RCSP mechanism where each of the N multiplexed streams is regulated before it is considered for transmission.

Since scheduling conflicts will arise when more than one regulated applications are present, a scheduler needs to be employed to resolve these conflicts. A consequence of the scheduling conflicts is that the tagged traffic stream at the output of the scheduler will be a distorted version of the target stream enforced at the output of the regulator. For instance, although a minimum spacing between consecutive tagged cells is enforced at the output of the tagged regulator in Fig. 1, this does not hold true for the tagged stream at the output of the scheduler. Clustering is generated due to an increased arrival rate to the scheduler in the immediate past which has pushed back (delayed) earlier tagged cells. Due to the latter, some spreading followed by some clustering of tagged cells is expected to be observed at the output of the scheduler.

The tagged cell spreading mentioned above can be reduced by monitoring the scheduler and releasing a tagged cell before its eligibility time[1] not allowing for the above conditions to be met. The dynamic-R&S scheme proposed below attempts to provide for a smoother tagged traffic at the output of the scheduler based on this idea.

Although less commonly stated, another objective in regulating real-time traffic streams within the network is to control (limit) the amount of bandwidth that is demanded by traffic streams. The spreading indicated above represents an instantaneous reduction in the bandwidth allocated to the tagged traffic stream, as measured at the output of the scheduler. In the context of the bandwidth availability to the tagged traffic stream,

[1]Here defined as T time units following the previous tagged cell release, if a minimum spacing of T is targeted.

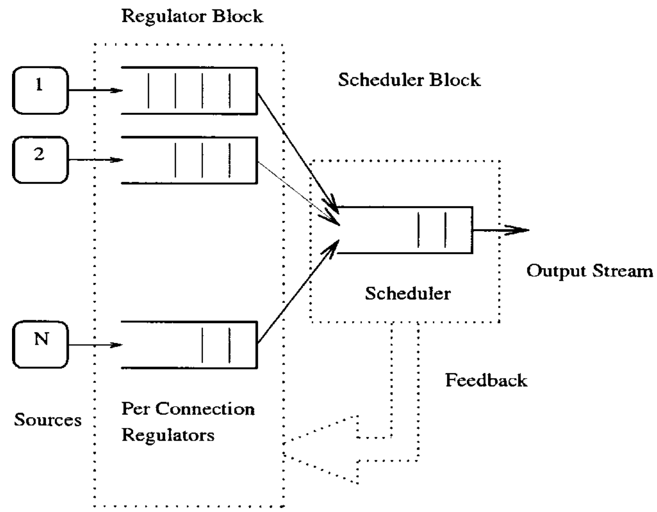

Fig. 1. The dynamic-R&S (including the arrow) and the static-R&S (excluding the arrow) systems.

the dynamic-R&S scheme proposed below may be viewed as attempting to provide for a constant bandwidth availability to the tagged traffic stream at periods of excessive total bandwidth demand from the coexisting applications.

A simple architecture of the switch for the illustration of the proposed policy is shown in Fig. 1. Similarly to the architecture proposed in [8], each of the N supported real-time applications is regulated at a logically dedicated regulator before it is delivered to the scheduler. In the present work, a simple FCFS scheduling policy is being considered. This scheduling mechanism is the simplest possible, reducing the scheduling complexity to single queue buffering.

Under the dynamic-R&S policy proposed below, the regulation process is modulated by some scheduler status information. Unlike in past work in the area, appropriate information regarding the status of the scheduler (FCFS queue) is fed back to the regulators, as indicated in Fig. 1 with the feedback arrow.

A. The Tagged Cell Release Policy: Dynamic-R&S Scheme

Let t_k denote the time slot at which the kth tagged cell is released from the tagged regulator. A slot is the service time of a single cell. Cells are of fixed length, and therefore time is measured in slots. The scheduler queue gets drained by one cell every time slot, given a nonempty queue. Otherwise, the slot expires unutilized. Let Q_k^r denote the queue occupancy at the regulator upon (following) the release of the kth cell. Let $t_k + B_k$ $(B_k \geq 1)$ denote the time slot at which the cumulative number of nontagged arrivals (releases) to the scheduler following t_k exceeds $T-2$ for the first time. Let a superscript d (s) indicate a quantity associated with the dynamic-R&S (static-R&S) policy, and let

$$W_k = \min\{B_k, T\} \text{ or } W = \min\{B, T\} \qquad (1)$$

where the last expression involves the generic random variables W and B. The $(k+1)$ tagged cell release time t_{k+1} is given by

$$t_{k+1} = t_k + W_k + H_k^d * 1_{\{Q_k^r=0, A_r^{d,W_k}=0\}} \qquad (2)$$

where H_k^d denotes the time interval between $\overline{t_k} = t_k + W_k$ and the first tagged cell arrival following $\overline{t_k}$; T is a constant positive integer; $A_r^{d,j}$ is the number of cell arrivals to the dynamic regulator over j slots; $1_{\{\text{expression}\}}$ is the indicator function, which is equal to one if the expression is true and zero otherwise. Fig. 2 illustrates the events associated with these definitions.

If T is equal to the minimum spacing among consecutive tagged cell releases from the regulator in the RCSP scheme [8], then it is easy to see that the above release policy will accelerate the tagged cell releases from the regulator at times when a minimum spacing of T at the output of the scheduler would be violated. This acceleration occurs when $B_k < T$. It is expected that the tagged cell release acceleration will have a positive impact on the tagged cell delay jitter and availed bandwidth. To quantify such benefits, the static-R&S scheme is considered in parallel in the rest of this paper. As described below, its tagged cell release policy is not modulated by any scheduler status information. A simple FCFS scheduler is also considered.

B. The Tagged Cell Release Policy: Static-R&S Scheme

By employing the definitions presented above and replacing $W_k = \min\{B_k, T\}$ by T, the $(k+1)$st tagged cell release time t_{k+1} is given by

$$t_{k+1} = t_k + T + H_k^s * 1_{\{Q_k^r=0, A_r^{s,T}=0\}} \tag{3}$$

where H_k^s denotes the time interval between $\overline{t_k} = t_k + T$ and the first tagged cell arrival following $\overline{t_k}$; $A_r^{s,j}$ is the number of cell arrivals to the static regulator over j slots.

III. Study of Regulator Behavior

The behavior of the R&S schemes is evaluated by investigating their impact on a specific stream (tagged stream). The traffic at the output of the regulators associated with the remaining $N-1$ applications is aggregated and forms the background traffic, which competes with the tagged traffic for resources at the scheduler. Let A^k denote the number of background cells delivered to the scheduler over k consecutive slots; let A_i denote the number of background cells delivered to the scheduler in the ith slot (note that $A^k = \Sigma_{i=1}^k A_i$). Since the input process to the scheduler is the output process from the regulator, it is important that the latter be investigated to both gain insight into the combined system (regulator plus scheduler) behavior and evaluate the output process at the scheduler.

The basic operational difference between the dynamic-R&S and static-R&S schemes is captured by the tagged cell interdeparture process from the regulator $\{V_k\}_{k\geq 1}$, where $V_k = t_{k+1} - t_k$. In view of (2) and (3), it is easy to establish that the evolution of the tagged cell process $\{V_k\}_{k\geq 1}$ is described by

$$V_k^s = T + H_k^s * 1_{\{Q_k^r=0, A_r^{s,T}=0\}} \quad \text{(static-R\&S)} \tag{4}$$

and

$$V_k^d = W_k + H_k^d * 1_{\{Q_k^r=0, A_r^{d,W_k}=0\}} \quad \text{(dynamic-R\&S)}. \tag{5}$$

In order to decouple the intrinsic behavior—to be investigated in this paper—of the R&S schemes from the source load, the heavy traffic source assumption will be made throughout this paper. This assumption is consistent with standard ones made in order to determine the throughput capabilities of a scheme as well as the throughput fluctuations (jitter), without the noise introduced by source inactivity periods. Under the heavy traffic source assumption, the regulator queue is considered nonempty, which implies that the indicator function in (4) and (5) is always zero. Thus

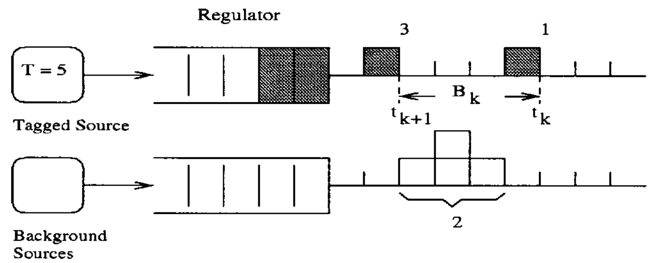

Fig. 2. Sequence of events describing the dynamic R&S release policy: 1) tagged cell k departs from the scheduler; 2) background exceeds $T-2$ cells; and 3) tagged cell $k+1$ gets released.

$$V_k^s = T \quad \text{(static-R\&S)} \tag{6}$$

and

$$V_k^d = W_k \quad \text{(dynamic-R\&S)}. \tag{7}$$

The following proposition describes the generic random variable W.

Proposition 1: The probability mass function of W and its mean are given by:

$$\Pr\{W = j\} = \tilde{F}_{j-1}(T-2) - \tilde{F}_j(T-2), \qquad 1 \leq j \leq T \tag{8}$$

and

$$E\{W\} = \sum_{j=1}^{T} \tilde{F}_{j-1}(T-2) \tag{9}$$

where $\tilde{F}_0(T-2) \triangleq 1$, $\tilde{F}_T(T-2) \triangleq 0$, and $\tilde{F}_j(T-2) \triangleq F_j(T-2)$, $1 \leq j < T$, where $F_j(\cdot)$ denotes the j-fold convolution of the probability distribution function of A^1. □

The proof of this proposition along with the proofs of other propositions and corollaries that follow may be found in [12].

Proposition 1 describes the regulator interdeparture process of the dynamic-R&S scheme in terms of the first-order probability mass function and the first moment. This process will be employed in the study of the scheduler under the dynamic-R&S scheme. Its description is also employed in the following comparative study of the two policies.

Proposition 2: The maximum throughput (output) rate of the tagged regulator under the two policies is given by

$$R_{\max}^s = \frac{1}{T} \tag{10}$$

and

$$R_{\max}^d = \frac{1}{E\{W\}} = \frac{1}{\sum_{j=1}^{T} \tilde{F}_{j-1}(T-2)}. \tag{11}$$

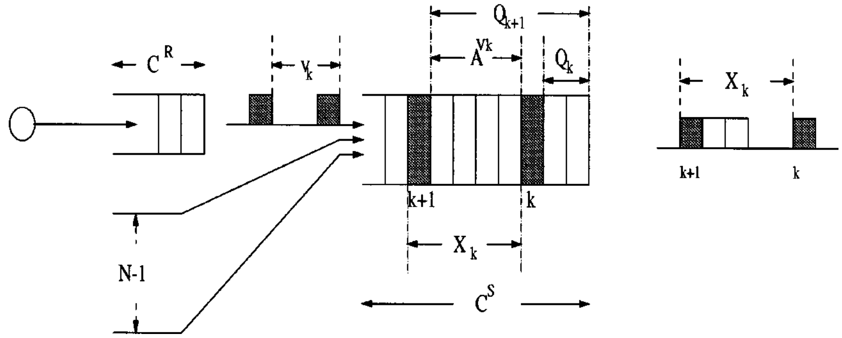

Fig. 3. Queueing model for the tagged traffic study at the scheduler.

Notice that $R^d_{\max} \geq R^s_{\max}$ with equality only when the background traffic process can never deliver more than $T-2$ cells over $T-1$ consecutive slots (typically, a zero probability event). □

The above discussion establishes that the dynamic-R&S scheme will respond to a sudden increase of the background load[2] by increasing its rate above the targeted rate of $1/T$, in an effort to ensure that the targeted tagged rate at the output of the scheduler is achieved. The impact of such a reaction (which is not possible under the static-R&S scheme) on the scheduler output process is investigated in the next section.

IV. Study of Scheduler Behavior Under Underload Conditions

In this section, the tagged cell interdeparture process from the scheduler is derived in order to evaluate the throughput/jitter properties of the R&S schemes. Since the deliverability of the QoS of the tagged application (which is considered to be a measure of the induced throughput/jitter) is expected to be decreased under high load conditions at the scheduler, the latter will be studied under high load conditions both below (underload) and above (overload) the scheduler capacity. Although the call admission control function will attempt to minimize the occurrence of temporary overload at the network nodes, due to traffic burstiness and the desire to achieve high resource utilization through statistical multiplexing, it is expected that temporary overload will be unavoidable. It is under such conditions that traffic management mechanisms should not only not collapse but minimize the impact of the overload on the QoS. In order to evaluate the tagged output process at the scheduler as shaped by the R&S policy as opposed to buffer overflows, sufficiently large buffer will be assumed to be available at the scheduler throughout this paper. To avoid unnecessary complications as well as evaluate the performance of the policies under nonidling environment—in which case they differ—the tagged source heavy traffic source assumption is maintained throughout the analysis.

Under underload conditions $(\rho < 1)$, the scheduler queue is stable, and since no cell overflow is possible, the following relationships between the maximum regulator $(R_{\max})$ and scheduler $(S_{\max})$ throughputs hold:

$$R^d_{\max} = S^d_{\max} \quad \text{and} \quad R^s_{\max} = S^s_{\max}. \tag{12}$$

In view of the above and Proposition 2, the following corollary is self-evident.

Corollary 1: $S^d_{\max} \geq S^s_{\max}$, where $S_{\max}$ denotes the maximum tagged cell output rate from the scheduler, or the maximum tagged cell throughput.

Corollary 1 implies that potentially higher throughput will be achieved by the tagged application under the dynamic-R&S policy compared to that under the static-R&S policy. This increase in throughput is due to a regulator tagged interdeparture interval W less than T under the dynamic-R&S policy, occurring only if the tagged cell interdeparture from the scheduler is to exceed the tagged value of T. Since the R&S policy attempts to minimize the variation of the target spacing T between consecutive tagged cell departures (or minimize the deviation of the instantaneous tagged cell output rate from $1/T$), the deviation of the tagged cell interdeparture interval from the target T will be employed as a jitter metric. The precise study of the tagged interdeparture process at the output of the scheduler—denoted by $\{X_k\}_{k\geq 1}$—is presented in the following subsections under both policies.

A. Tagged Cell Interdeparture Under the Dynamic-R&S Policy

A queueing model for the system of the tagged traffic regulator and scheduler is shown in Fig. 3; the regulator (C^R) and scheduler (C^S) buffer capacities are assumed to be theoretically infinite, as mentioned earlier. Also as indicated earlier, the heavy traffic assumption implies the regulator will not starve, and thus V_k depends only on $W_k = \min\{B_k, T\}$, where B_k is the shortest time interval (in slots), initiated upon the release of the kth tagged cell and terminated at the slot in which the cumulative nontagged (or background) arrivals over this interval exceed $T-2$. As implied in Fig. 3, $N-1$ traffic streams—potentially controlled by a similar R&S scheme—are assumed to compete for the common, slotted transmission resource. For analysis tractability, the cumulative arrivals from the $N-1$ coexisting traffic streams—forming the background traffic—are assumed to be independent (per slot) and identically distributed, represented by the random variable A^1 (see above); the maximum number of background arrivals per slot is equal to $N-1$.

[2]This would result in an instantaneous reduction of the tagged throughput at the output of the scheduler.

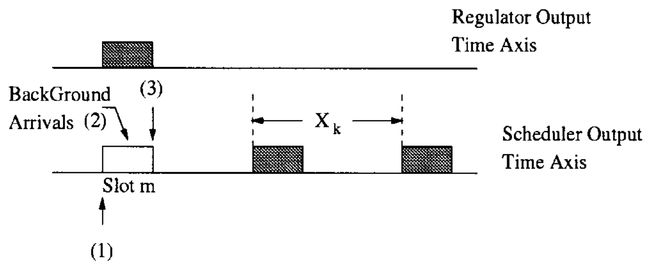

Fig. 4. Sequence of events in the scheduler output slot m: 1) cell departure; 2) background cell arrivals; and 3) tagged cell arrival.

Fig. 4 illustrates the sequence of events associated with a slot as considered in this study. It is understood that cell arrivals will rarely occur exactly on the slot boundaries. Cells arriving during a time slot will be considered for service during the next slot according to following convention. A cell departure (if any) is assumed to occur first, followed by the cumulative background arrivals over this slot (if any) and then the tagged cell arrival (if any). X_k-1 slots are available to the background traffic between two consecutive tagged cell departures with interdeparture interval equal to X_k. The scheduler interdeparture time between tagged cells k and $k+1$, X_k, can be determined based on the scheduler queue occupancy Q_k found upon arrival of the tagged cell k to the scheduler queue (not counting itself). This is described next. Later, the stationary probabilities of Q_k are derived to complete the calculation of the probability distribution of X_k. A rarely encountered peculiarity of the scheduler queue is that tagged cell arrivals depend on the queue build up, or the queue's arrival process depends on its occupancy process. Specifically, tagged cell arrival $k+1$ occurs T slots following the arrival of tagged cell k if and only if the cumulative background arrivals following ("behind") the tagged cell k arrival have not exceeded $T-1$ at any earlier slot. Otherwise, tagged cell $k+1$ will arrive (be released from the regulator) at that earlier slot. The tagged cell k will be served in (during) the (Q_k+1) slot following its release from the regulator. The following proposition provides for the conditional value of X_k given Q_k.

Proposition 3: The conditional probability of X_k given Q_k, $\Pr\{X_k=T+l/Q_k=i\}$ for $-(T-1)\le l\le (N-2)$, is given by the following expressions.

Case I: $l\ge 0$ (no clustering)

$$\begin{aligned}\Pr\{X_k=T+l/Q_k=i\} &= \sum_{m=1}^{\min\{i+1,T\}} \Pr\{A^m=T+l-1, A^{m-1}\le T-2\}\\ &\quad + 1_{\{i\le T-2\}} * \sum_{m=0}^{T-2} \Pr\{A^{i+1}=m, X_k=T+l\}.\end{aligned} \tag{13}$$

Case II: $l<0$ (clustering)

$$\begin{aligned}\Pr\{X_k=T+l/Q_k=i\} &= 1_{\{i\ge T-1\}} * \Pr\{A^T=T+l-1\}\\ &\quad + 1_{\{i\le T-2\}} * \sum_{m=0}^{T-2} \Pr\{A^{i+1}=m, X_k=T+l\}\end{aligned} \tag{14}$$

where the probabilities involved in the above expressions are derived and evaluated in the proof of this proposition. □

The transition probabilities of the scheduler occupancy process $\{Q_k\}_{k\ge 1}$ are given in the following proposition.

Proposition 4: The transition probabilities of the Markov process $\{Q_k\}$, embedded upon tagged cell arrival times, are given by the following expressions.

Case I: $i\ge T-1$

$$\begin{aligned}\Pr\{Q_{k+1}=j/Q_k=i\} &= \Pr\{A^1=j-i\} * 1_{\{T-1\le j-i\le N-1\}\cap\{T-1\le N-1\}}\\ &\quad + \sum_{m=2}^{T-1} \Pr\{A^m=j-i-1+m, A^{m-1}\le T-2\}\\ &\quad * 1_{\{T-m\le j-i\le T-2+N-m\}\cap\{T-1\le m*(N-1)\}}\\ &\quad + \Pr\{A^T=j-i-1+T\} * 1_{\{-(T-1)\le \supset -i\le N-2\}}.\end{aligned}$$

Case II: $i\le T-2$

$$\begin{aligned}\Pr\{Q_{k+1}=j/Q_k=i\} &= \Pr\{A^1=j-i\} * 1_{\{T-1\le j-i\le N-1\}\cap\{T-1\le N-1\}}\\ &\quad + \sum_{m=2}^{i+1} \Pr\{A^m=j-i-1+m, A^{m-1}\le T-2\}\\ &\quad * 1_{\{T-m\le j-i\le T-2+N-m\}\cap\{T-1\le m*(N-1)\}}\\ &\quad + \sum_{m=0}^{T-2}\sum_{n=1}^{T-i-2}\sum_{a_n=T-1-m}^{T-3+N-m} \Pr\{A^{i+1}=m, A^i\le T-2\}\\ &\quad * \Pr^n\{\hat{A}(n)=a_n, \hat{Q}(n)=j/\hat{A}(0)=0, \hat{Q}(0)=m\}\\ &\quad + \sum_{m=0}^{T-2}\sum_{a=0}^{T-3+N-m} \Pr\{A^{i+1}=m, A^i\le T-2\}\\ &\quad * \Pr^{T-i-1}\{\hat{A}(T-i-1)=a, \hat{Q}(T-i-1)=j/\\ &\quad \cdot /\hat{A}(0)=0, \hat{Q}(0)=m\}.\end{aligned}$$

Last, the tagged cell interdeparture probability distribution can be obtained from the conditional ones (Proposition 3) and the stationary probabilities of $\{Q_k\}$, $\pi_q(i)$, derived by employing the transition probabilities given in Proposition 4. Thus

$$\begin{aligned}\Pr\{X_k=T+l\} &= \sum_{i=0}^{\infty} \Pr\{X_k=T+l/Q_k=i\} * \pi_q(i)\\ &= \sum_{i=0}^{T-2} \Pr\{X_k=T+l/Q_k=i\} * \pi_q(i)\\ &\quad + \left[1-\sum_{i=0}^{T-2}\pi_q(i)\right] * \Pr\{X_k=T+l/Q_k\ge T-1\}\end{aligned} \tag{15}$$

Since the conditional interdeparture given $Q_k=i$ for $i\ge T-1$ is constant, independent from i (see Proposition 3), (15) suggests that since only the stationary probabilities $\pi_q(i)$ for $i\le T-2$ are used, truncating $\{Q_k\}_{k\ge 1}$ to some state $T-2+L$, for some large L, will cause negligible impact on the $\pi_q(i)$ for

$i \leq T-2$ since only values away from the truncation boundary are used in (15).

B. *Tagged Cell Interdeparture Under the Static-R&S Policy*

Under the static-R&S policy, no early tagged cell releases are allowed, and therefore the scheduler is fed by a periodic tagged traffic of period T (heavy traffic assumption at the regulator).

Proposition 5: The probability distribution of the jitter for the static-R&S policy is given by

$$\Pr\{X_k = T + l\} = \sum_{j=1}^{\infty} \Pr\{Q_{k+1} = i + l / Q_k = i\} * \pi_q(i). \quad (16)$$

□

The $\Pr\{Q_{k+1} = j/Q_k = i\}$ (transition probabilities) can be obtained as the T-step transition probabilities of a simple $M/D/1$ queue. By employing the transition probabilities, the stationary distribution $\pi_q(i)$ can be obtained.

V. Study of Scheduler Behavior Under Overload Conditions

Since the traffic streams that are traversing a single multiplexer can be of highly variable rates, it is possible—for a short period of time—to have an average arrival higher than one. If this condition were to persist, then the buffers would grow without bound and the system would be unstable. It is desirable to study the behavior of the policies under such extreme conditions. The case in which the scheduler queue is unstable is considered next. That is, $\rho > 1$. Although the scheduler queue capacity is again assumed to be infinite and, for that matter, no overflow will occur, the relationship in (12) will not hold since traffic will accumulate in the infinite queue. This ever increasing queue buildup will represent a difference between input and output traffic rate, making it hard to assess the precise throughput achieved by the tagged or background streams.

In view of the large buffer assumption at the scheduler, it is evident that the scheduler queue occupancy upon tagged cell arrival Q_k will always exceed $T-1$ under overload conditions $(\rho > 1)$. As a consequence—as will become evident below—the analysis of the scheduler tagged cell interdeparture process is simplified significantly. For this reason, both policies are treated concurrently.

The following proposition provides for the precise description of the tagged stream interdeparture distribution at the output of the scheduler under overload conditions.

Proposition 6: Under the overload conditions at the scheduler queue $(\rho > 1)$ and infinite scheduler buffer capacity, the distribution of the tagged cell interdeparture at the scheduler X_k under the dynamic-R&S policy is given by

$$\Pr\{X_k = l\} = \sum_{m=1}^{T} \Pr\{A^m = l-1, A^{m-1} \leq T-2\} * 1_{\{T \leq l \leq N+T-3\}} + \Pr\{A^T = l-1\} * 1_{\{1 \leq l \leq T-1\}}, \quad \text{for } 1 \leq l \leq N+T-3. \quad (17)$$

The distribution under the static-R&S policy is given by

$$\Pr\{X_k = l\} = \Pr\{A^T = l-1\}, \quad \text{for } 1 \leq l \leq (N-1)*T+1. \quad (18)$$

□

Proposition 6 can be employed in deriving the throughput of the tagged application. The following proposition establishes the better jitter and smoothness characteristics of the dynamic-R&S scheme, compared to those of the static-R&S scheme under overload conditions at the scheduler.

Proposition 7: Under overload conditions at the scheduler, the dynamic-R&S policy will potentially reduce the tagged cell spreading (while it will never increase it) compared to the static-R&S policy. That is

$$\Pr\{X_k^d > m\} \leq \Pr\{X_k^s > m\} \quad \text{for } m > T.$$

The tagged cell clustering will be identical under both policies, that is

$$\Pr\{X_k^d < m\} = \Pr\{X_k^s < m\}, \quad \text{for } 1 \leq m < T$$

and

$$\Pr\{X_k^d = T\} \geq \Pr\{X_k^s = T\}.$$

□

The following proposition establishes some insightful relationships between the moments of interdeparture and the background traffic process.

Proposition 8: Under overload conditions at the scheduler, the following can be shown:

$$\begin{aligned}
&\text{(a)} \quad E\{X_k^s\} = T * E\{A^1\} + 1 \\
&\qquad\qquad\quad = T * (N-1) * p_b + 1 \\
&\text{(b)} \quad VAR\{X_k^s\} = T * VAR\{A^1\} \\
&\qquad\qquad\quad = T * (N-1) * p_b * (1-p_b) \\
&\text{(c)} \quad E\{X_k^d\} = E\{W\} * E\{A^1\} + 1 \\
&\qquad\qquad\quad = E\{W\} * (N-1) * p_b + 1
\end{aligned}$$

where the last part of the above equations is derived for a binomial random variable A^1 with maximum value $N-1$ and success probability p_b; $VAR\{x\}$ denotes the variance of random variable x. □

VI. Numerical Results

In this section, some numerical results are presented to quantify the behavior of the dynamic-R&S and static-R&S schemes. The results are derived under heavy traffic source load at the regulator and both underload and overload conditions at the scheduler. Although the heavy traffic source and overload conditions are not the dominant ones in a well-designed system, they will be present if substantial statistical multiplexing gain is to be achieved. And while simple regulator schemes—such as the static-R&S one—may be adequate under nominal (underload) traffic conditions, it is important that their behavior under less frequent—but QoS compromising—overload conditions be investigated.

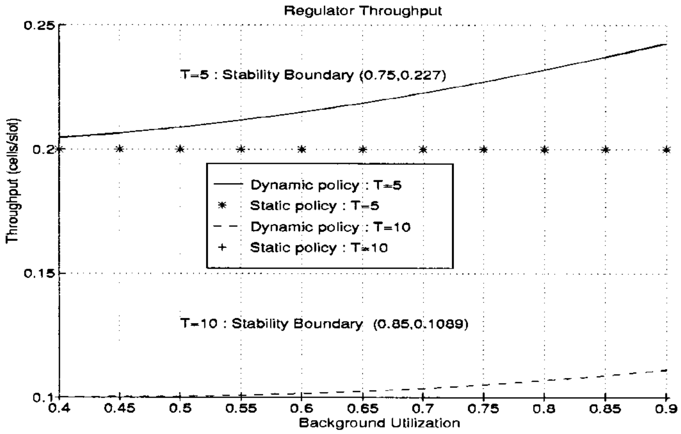

Fig. 5. Regulator throughput versus background utilization ρ_{back} ($\rho_{\text{back}} < 1$).

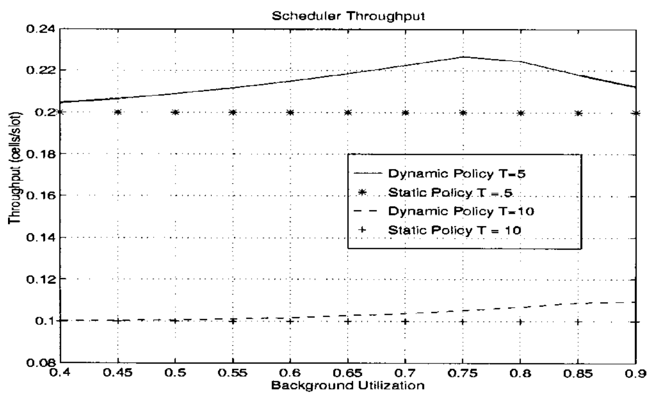

Fig. 6. Scheduler throughput versus utilization ρ ($\rho_{\text{back}} < 1$).

The results presented below have been obtained for two values of the target interdeparture time T—or desirable throughput $1/T$—equal to $T = 5$ and $T = 10$. The background traffic is modeled as an independent per slot, batch process with binomially distributed batch size of maximum value $N - 1 = 8$ and success probability p_b. The background load or utilization is denoted by ρ_{back}. Fig. 5 presents the regulator throughput versus ρ_{back} under both policies, obtained from Proposition 2. As ρ_{back} increases, the dynamic-R&S policy can detect the increased background intensity and release cells earlier, attempting to provide the targeted throughput $(1/T)$ and control jitter. As a consequence, the rate by which the packets leave the regulator increases as the background intensity increases, and it will reach a maximum of one if at least $T-1$ background cells are delivered to the scheduler in each slot (high overload at the scheduler).

The scheduler throughput $S_{\max}$ versus ρ_{back} is shown in Fig. 6 for both policies under underload conditions ($\rho < 1$). $S_{\max}$ is calculated as $1/E\{X_k\}$; the probability mass function (PMF) of X_k is calculated from (15) for the dynamic policy, and from (16) for the static policy. As expected from (12), $S^d_{\max} = R^d_{\max}$ under underload conditions and infinite buffer capacity at the scheduler. Although $R^d_{\max}$ increases to one as ρ_{back} increases (as said earlier), $S^d_{\max}$ starts deviating from $R^d_{\max}$ and declines beyond some value of ρ_{back} equal to about 0.78 for $T = 5$. This is due to the fact that the scheduler

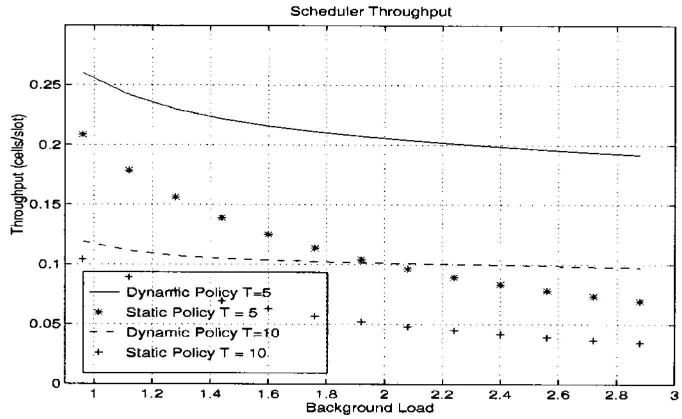

Fig. 7. Scheduler throughput versus background utilization ρ_{back}.

reaches an overload state ($\rho > 1$) and the dynamics change. Thus, only the results for $R_{\max} + \rho_{\text{back}} < 1$ are relevant. The scheduler study under overload should be employed for the derivation of results for $\rho > 1$. Results for $S_{\max}$ versus ρ_{back} are shown in Fig. 7 for ρ_{back} such that the scheduler is in overload state ($\rho > 1$). The results are calculated from the PMF of X_k derived in (18) and (17). The improved throughput characteristics of the dynamic-R&S scheme can be clearly observed. The higher than the targeted throughput for low overload conditions is in accordance with expectations based on the increased regulation throughput and the heavy traffic assumption for the tagged source. As long as $\rho < 1$, the regulator throughput determines the throughput at the scheduler as well, as explained earlier. When $\rho > 1$, some of the regulator traffic is "absorbed" by the infinite buffer built up at the scheduler, and a throughput reduction is observed for this reason. Nevertheless, by reducing W (see (1)), the dynamic-R&S scheme is capable of providing the targeted throughput under severe overload conditions. In the limiting case of very large overload, $W \to 1$ and the tagged throughput reduction below the targeted value is observed, induced by the per slot background batch size. It should be noted that under the static-R&S scheme, the tagged throughput falls dramatically even under low overload conditions. This reduction is directly related to the cumulative over T slots background arrivals, as opposed to that over W slots $(W \to 1)$ under the dynamic-R&S scheme.

Fig. 8 presents results for the variance of the tagged cell interdeparture process induced by the two policies versus ρ_{back} and for underload conditions at the scheduler ($\rho < 1$). Again, only the results for $R_{\max} + \rho_{\text{back}} < 1$ are relevant. It is clear that the dynamic-R&S policy provides for a less variable interdeparture process than the static-R&S one.

Similar results for $\rho > 1$ are presented in Fig. 9. These results have been derived by employing the PMF of X_k derived in Proposition 6, for various values of ρ_{back}. In view of the linear relationship between $VAR(X_k)$ and $VAR(A^1)$ under the static-R&S scheme (Proposition 8) the increasing behavior of $VAR(X_k)$ as ρ_{back} increases is expected, and it is observed in Fig. 9. The results under the dynamic-R&S scheme are more difficult to interpret. For low overload conditions, $VAR(X_k)$ decreases until $\rho_{\text{back}} = 1.44$ (fourth point on the plot) and then increases slightly. X_k depends solely on the

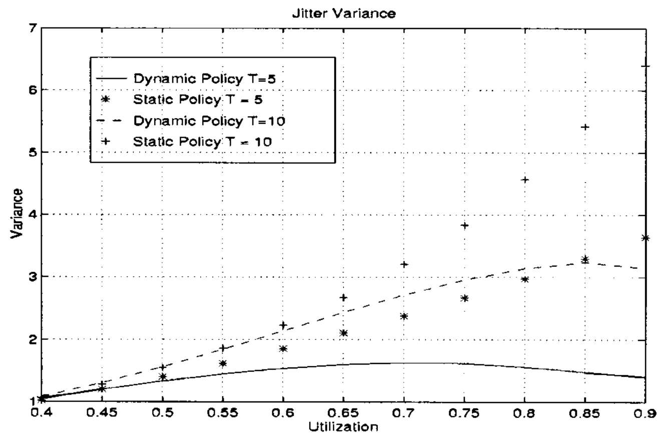

Fig. 8. Variance versus background utilization ρ_{back} ($\rho_{\text{back}} < 1$).

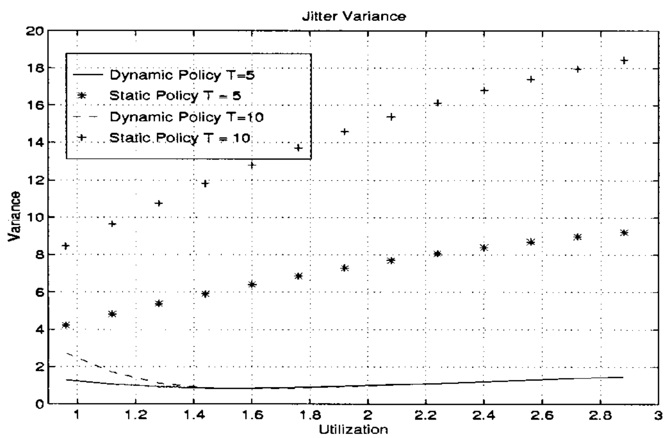

Fig. 9. Variance versus background utilization ρ_{back} ($\rho_{\text{back}} >, 1$). $P_b = 0.12 + (k-1) * 0.02$ where k the point of interest ($k = 1, 2, \ldots, 13$).

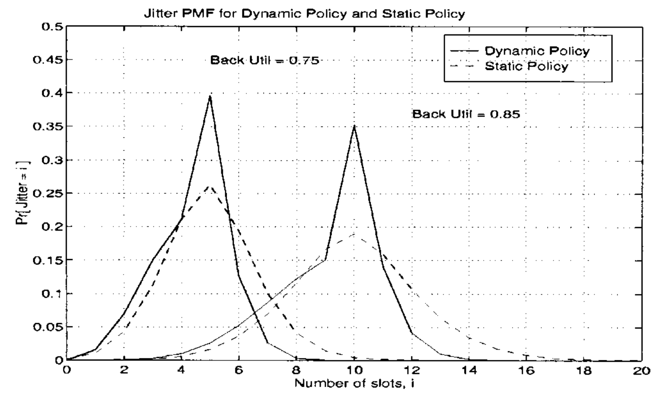

Fig. 10. High link utilization: jitter performance improves for the dynamic policy ($\rho_{\text{back}} < 1$).

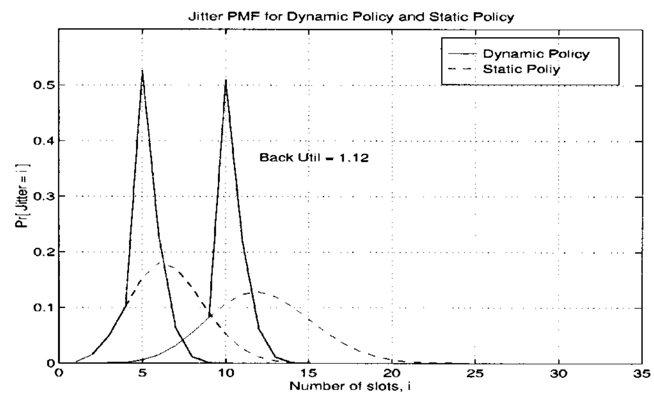

Fig. 11. Jitter distributions under moderate overload ($\rho_{\text{back}} > 1$).

background accumulation over W slots (between consecutive tagged cell releases). Therefore, as ρ_{back} increases, the condition $A^W \geq T - 1$ is expected to be met in fewer slots, and therefore a cell would be released earlier. This implies that a decreasing number of batches would interfere with X_k, reducing $VAR(X_k)$. As ρ_{back} increases the number of batches that interfere with X_k under the dynamic-R&S policy reduces to one; beyond that point, the increased value of $VAR(X_k)$ is due to the increase in $VAR(A^1)$. The jitter (X_k) PMF under both policies is plotted in Fig. 10 for $\rho_{\text{back}} = 0.75$, $\rho_{\text{back}} = 0.85$ (Back Util), and $T = 5$ and $T = 10$. The PMF of X_k becomes quite distinct for the two policies. Tagged cell clustering $(X_k < T)$ is seen to slightly increase under the dynamic-R&S policy while spreading $(X_k > T)$ is substantially reduced and more probability mass is concentrated around T. While the spreading reduction under the dynamic-R&S policy is expected, the slight increase in the clustering is less obvious. It may be attributed to the higher probability that the scheduler queue is nonempty under the dynamic-R&S policy, due to the higher scheduler load $(S_{\max})$ resulting from a higher regulator throughput $(R_{\max})$. This is also discussed below.

The traffic smoothness characteristics of the two schemes under overload traffic conditions at the scheduler can be observed in Fig. 11, where the jitter distribution (or scheduler interdeparture distribution) is plotted for $\rho_{\text{back}} = 1.12$. The jitter probability mass function is highly contained around the target value T under the dynamic-R&S scheme, and it is spread over wide range of values under the static-R&S scheme.

The good jitter characteristics of the dynamic-R&S scheme in terms of reduced cell spreading are clearly observed. Excessive spreading—occurring under network congestion (scheduler overload)—may compromise the QoS of a real-time application by causing starvation at the end user. It is evident that the starvation probability can be substantially lower under the dynamic-R&S scheme.

For more results and analysis on the traffic smoothness properties of the two policies, please refer to [12].

VII. Simulation Results

In this section, a system in which more than one sources are controlled by the R&S policies is considered and is studied using OPNET. The objective here is to investigate the behavior of the two policies in the presence of real background traffic, as generated by multiple sources controlled by these policies.

A system with $N = 7$ ON_OFF Markov sources was simulated. Reference [12] contains details about the source parameters and other simulation parameters. The actual cell interdeparture times (X_k) from the scheduler were recorded, after filtering out the gaps caused by the source's OFF periods, and a vector $\overline{X} = [X_1 \ X_2 \ \ldots \ X_k \ \ldots]$ was created for each of the sources. The empirical interdeparture PMF and throughput for each source was obtained from the samples in $\overline{X}$. Since $\overline{X}$

TABLE I
SIMULATION PARAMETERS FOR THE FIRST SIMULATION SCENARIO, WHERE X IS $X_{\min}$, $SMG = \Sigma_i\ 1/X_{\min,i} = 3.6$, $\rho = \Sigma_i\ \lambda_{\text{ave},i} = 0.9505$, VAR_d IS $VAR[X_k^d]$ AND VAR_s IS $VAR[X_k^s]$

src	X	λ_{ave}	S_a^d	S_a^s	VAR_d	VAR_s
0	1	0.247	0.449	0.455	3.260	3.023
1	1	0.195	0.396	0.402	4.865	4.604
2	1	0.200	0.402	0.409	4.625	4.392
3	10	0.043	0.108	0.096	4.624	9.336
4	10	0.042	0.106	0.094	4.029	8.886
5	5	0.111	0.203	0.182	3.942	5.185
6	5	0.112	0.207	0.186	3.306	4.505

does not contain the OFF periods, the calculated throughput will be higher than the actual average throughput. This calculated throughput is called the active throughput S_a. This filtering of the data allows for capturing the effect that the policies have on the sources while they are active, without being obscured by the OFF periods where the policies are ineffective.

For each source, the value of $X_{\min}$ and the measured average source rate λ_{ave} are shown, along with the measured active throughput S_a^d and variance $VAR[X_k^d]$ under the dynamic-R&S and S_a^s and variance $VAR[X_k^s]$ under the static-R&S policies.

Table I shows the parameters for the 1st simulation scenario. The Statistical Multiplexing Gain (SMG) is equal to 3.6 in this case. The system utilization, ρ, is equal to 0.9505, which should be less than 1 for a stable system. The sources with $X_{\min} = 1$ are practically unregulated since they are allowed to release their cells to the scheduler as soon as they are generated. Such sources could be ones without jitter constraints. Sources with $X_{\min} > 1$ are the ones targeted for regulation; for these sources, $X_{\min}$ is set to $\lfloor 1/P_{\text{arr}} \rfloor$, where P_{arr} is the cell generation rate while the source is on the ON state.

The positive impact of the dynamic-R&S policy on the regulated sources $(src_3, src_4, src_5, src_6)$ is clearly observed in Table I: the variance of the interdeparture process under the dynamic-R&S scheme $(VAR[X_k^d])$ is substantially smaller than the variance under the static-R&S scheme $(VAR[X_k^s])$. This quantifies the ability of the dynamic-R&S scheme to control the delay jitter better than the static-R&S scheme, for all of the regulated sources. As a result, under the dynamic-R&S scheme, S_a^d reaches the target value $(1/X_{\min})$ and exceeds it slightly, due to some residual clustering which is present under both policies. On the other hand, under the static-R&S scheme S_a^s is lower than the target value due to the static nature of the policy. The dynamic-R&S scheme manages to serve the regulated sources with a less variable service rate, which reaches the target peak service rate. As a result, the traffic is better shaped at the output of the scheduler, producing the desired performance.

The unregulated sources (src_0, src_1, src_2) experience a slightly higher $VAR[X_k^d]$. Since the target $X_{\min}$ for these sources is one, the two policies are basically ineffective and the variance and active throughput are shaped solely by the background interfering process seen by each of these sources during each slot. It is expected that under the dynamic-R&S policy more cells will be released to the scheduler per slot. Therefore, more background traffic interferes with the unregulated sources.

The above results can be viewed graphically in terms of the sample empirical PMF's in Figs. 12–15. The PMF at the target value is substantially higher under the dynamic-R&S scheme.

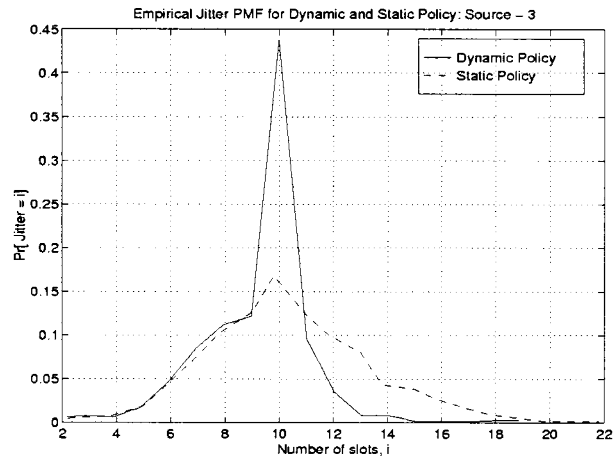

Fig. 12. Empirical PMF's—first simulation senarion.

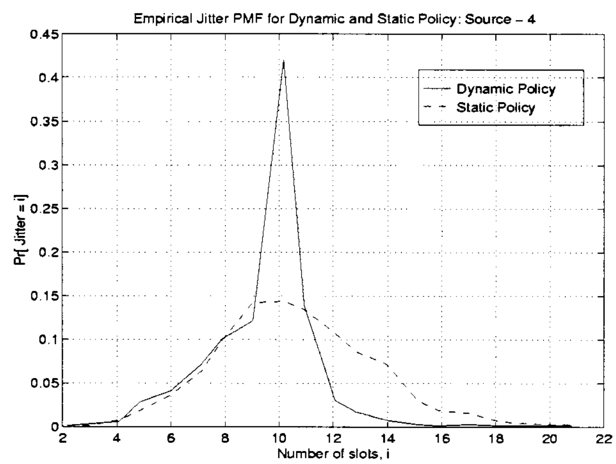

Fig. 13. Empirical PMF's—first simulation senarion.

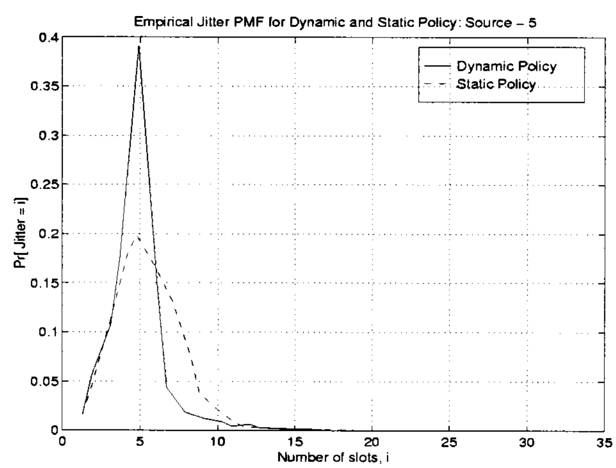

Fig. 14. Empirical PMF's—first simulation senarion.

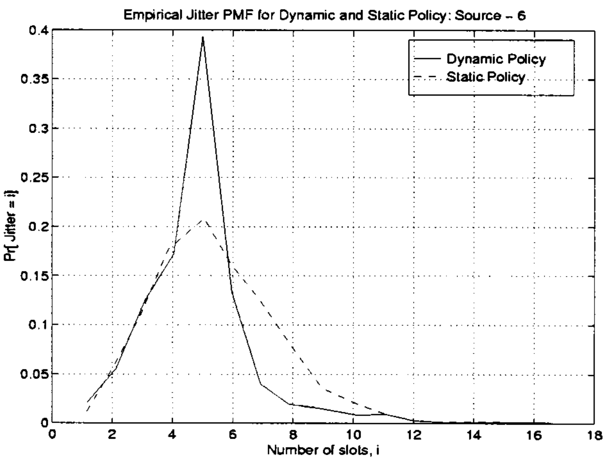

Fig. 15. Empirical PMF's—first simulation senarion.

TABLE II
SIMULATION PARAMETERS FOR THE SECOND SIMULATION SCENARIO, WHERE X IS $X_{\min}$, $SMG = \Sigma_i\ 1/X_{\min,i} = 3.6$, $\rho = \Sigma_i\ \lambda_{\text{ave},i} = 1.4905$, VAR_d IS $VAR[X_k^d]$, AND VAR_s IS $VAR[X_k^s]$

src	X	λ_{ave}	S_a^d	S_a^s	VAR_d	VAR_s
0	1	0.398	0.430	0.435	3.123	2.949
1	1	0.384	0.429	0.431	3.133	2.980
2	1	0.400	0.433	0.437	3.027	2.913
3	10	0.043	0.093	0.065	8.262	48.03
4	10	0.042	0.091	0.064	8.987	51.96
5	5	0.111	0.158	0.124	11.99	20.32
6	5	0.112	0.163	0.126	9.826	17.71

This may be attributed to the reduced spreading. Everytime a cell is released earlier than the release time under the static-R&S policy, a potential spreading is avoided. The empirical PMF's for $X_k < X_{\min}$ almost coincide under the two policies, implying that the two policies generate the same amount of clustering. The empirical PMF's of src_0, src_1, src_2 are almost identical under the two policies, and they are omitted since they do not provide any insight.

In order to study the system under extreme overload, the average rate of the three unregulated sources was increased (Table II). This could reflect a scenario according to which the unregulated sources start to misbehave and become bandwidth greedy. The utilization of the system was brought up to 1.4905 for some time; then the sources were turned off and the scheduler was served until it became empty. This way, the behavior of the two policies under temporary overload conditions could be studied.

Table II shows results for the variance of the interdeparture process and active throughput under this scenario. The variance of the regulated sources under the static-R&S policy increases dramatically. This is not the case under the dynamic-R&S policy; which manages to keep the variance substantially lower. Even though under extreme overload the analytical results predict that S_a^d should be fairly close to the target value $1/X_{\min}$ (Fig. 7), it is not seen here. The latter is due to the heavy traffic assumption made in the analytical study, under which the regulator never empties and the maximum effect of the policy can be revealed. In the simulations, the regulator can be empty when the conditions for a cell release are met. The eligible cell will be released in a later slot as soon as it arrives, allowing the misbehaving sources to secure more bandwidth.

VIII. CONCLUSION

In this paper; dynamic regulation and scheduling scheme has been proposed and studied through analysis and simulation. The scheme inspired the formulation and solution of a challenging queueing problem, where the arrival process to the scheduler was dependent on the scheduler queue occupancy. The proposed policy was studied under conditions of high link utilization as well as temporary overload conditions. Both simulation and analytical studied focused on a single node scenario. The dynamic R&S policy was compared against a static counterpart of comparable complexity and goals. The analytical and simulation results clearly showed that the dynamic-R&S scheme outperformes its static counterpart. It has been shown that the dynamic-R&S scheme can provide substantially better jitter control and achieve higher statistical multiplexing gain than the static-R&S scheme. The improved jitter characteristics of the dynamic R&S scheme yielded a less variable peak service rate seen by the regulated sources. A corrolary of the jitter reduction properties is the "fair" allocation of bandwidth between competing sources. Further study may focus on loosening some of the assumptions that were maintained throughout the analysis, on an end-to-end study and on a buffer management scheme so that the infinite buffer assumptions can be relaxed.

REFERENCES

[1] D. Ferrari and D. Verma, "A scheme for real time channel establishment in wide area networks," *IEEE J. Select. Areas Commun.*, vol. 8, pp. 368–379, Apr. 1990.
[2] D. Ferrari, "Delay jitter control scheme for packaging switching internetwork," *Comput. Commun.*, vol. 15, no. 6, pp. 367–373, July 1992.
[3] D. C. Verma, H. Zhang, and D. Ferrari, "Delay jitter control for real-time communication in a packet switching network," in *Proc. Tricomm '91*, Chapel Hill, NC, Apr. 1991, pp. 35–46.
[4] C. R. Kalmanek, H. Kanakia, and S. Keshav, "Rate controlled servers for very high speed networks," in *Proc. IEEE Global Telecommunication Conf.*, San Diego, CA, Dec. 1990, pp. 300.3.1–300.3.9.
[5] S. J. Golestani, "Congestion-free communication in high-speed packet networks," *IEEE Trans. Networking*, vol. 39, pp. 1801–1812, Dec. 1991.
[6] A. Demers, S. Keshav, and S. Shenker, "Analysis and simulation of a fair queueing algorithm," in *Proc. ACM SIGCOMM '89*, Oct. 1989, pp. 1–12.
[7] A. K. Parekh and R. G. Galager, "A generalized processor sharing approach to flow control in integrated services networks: The single-node case," *IEEE/ACM Trans. Networking*, vol. 1, pp. 344–357, June 1993.
[8] H. Zhang and D. Ferrari, "Rate-controlled static-priority queueing," in *Proc. IEEE INFOCOM '93*, Sept. 1993, pp. 227–236.
[9] N. R. Figueira and J. Pasquale, "Leave in time: A new service discipline for real time communications in a packet switching network," in *Proc. ACM SIGCOMM '96*, Cambridge, MA, 1995, pp. 207–218.
[10] D. Saha, S. Mukherjee, and S. K. Tripathi, "Multi-rate traffic shaping and end-to-end performance guarantees in ATM Networks," in *Proc. Int. Conf. Network Protocols*, 1994, pp. 188–195.

[11] L. Zhang, "Virtual clock: A new traffic control algorithm for packet switching networks," in *Proc. ACM SIGCOMM '90*, Sept. 1990, pp. 19–29.

[12] S. Iatrou, "A dynamic regulation and scheduling scheme for real-time traffic management," M.S. thesis, Northeastern Univ., 1997.

Steve Iatrou received the B.Sc. degree from the University of Massachusetts, Lowell, and the M.Sc. degree from Northeastern University, Boston, MA, both in electrical engineering.

He is currently working for Andersen Consulting, Sophia Antipolis, France, in the Network Technology group.

Ioannis Stavrakakis (S'84–M'88–SM'93) received the diploma from the Aristotelian University of Thessaloniki, Greece, in 1983 and the Ph.D. degree from the University of Virginia, Charlottesville, in 1988, both in electrical engineering.

He was an Assistant Professor in computer science and electrical engineering, University of Vermont, during 1988–1994 and an Associate Professor of electrical and computer engineering, Northeastern University, Boston, MA, during 1994–1999;. He is currently and Associate Professor of Informatics, University of Athens, Greece. His teaching and research interests are focused on resource allocation protocols and traffic management for communication networks. His past research has been published in more than 90 scientific journals and conference proceedings. His research has been funded by NSF, DARPA, GTE, BBN and Motorola, as well as Greek and European Union funding agencies. He has served on NSF research proposal review panels and was involved in the organization of numerous conferences sponsored by IEEE, ACM, ITC, and IFIP societies. He is an Associate Editor of the *ACM/Baltzer Wireless Networks Journal*.

Prof. Stavrakakis is a member of the IEEE Technical Committee on Computer Communications (TCCC) and of IFIP WG6.3. He has served as an elected officer for TCCC, as a Co-organizer of the 1996 International Teletraffic Congress (ITC) Mini-Seminar, on "Performance Modeling and Design of Wireless/PCS Networks," as the Organizer of the 1999 IFIP WG6.3 workshop, and as a Technical Program Cochair for the IFIP Networking'2000 conference.