ATL: An Adaptive Transport Layer Suite for Next-Generation Wireless Internet

Özgür B. Akan, *Student Member, IEEE,* and Ian F. Akyildiz, *Fellow, IEEE*

***Abstract*—The next-generation wireless Internet (NGWI) is expected to provide a wide range of services including high-speed data and real-time multimedia to mobile users. To realize this expectation, a diverse set of challenges need to be addressed, which are posed by heterogeneous wireless networking environments within NGWI and the according application requirements. Furthermore, the architectural heterogeneities must be captured dynamically, while mobile users may roam during their connection duration. Current existing transport layer protocols have been developed for a specific network paradigm in mind, e.g., for wireless local area networks (WLANs), micro/macro wireless systems, or for satellite systems. Using these existing different transport layer protocols for NGWI to support global roaming of mobile users is not a practical solution due to processing and memory constraints of wireless terminals. Thus, there is a need for a unified adaptive transport layer protocol suite which can address the architectural heterogeneities for roaming mobile users and achieve the best performance for NGWI.**

In this paper, a unified adaptive transport layer (ATL) suite is introduced for NGWI which incorporates a new adaptive transport protocol (TCP-ATL) for reliable data transport and a new adaptive rate control protocol (RCP-ATL) for multimedia delivery in the NGWI. According to the requested service type, i.e., reliable data or multimedia, ATL selects the appropriate protocol. Both TCP-ATL and RCP-ATL, deploy a new adaptive congestion control method that dynamically adjusts the protocol configurations according to the current wireless network paradigms depending where the mobile user currently resides. Hence, the unified adaptive ATL protocol suite achieves high-throughput performance in all of underlying heterogeneous wireless architectures, i.e., WLANs, micro, macro, or satellite environments. Moreover, the developed adaptive congestion control explicitly takes fairness into consideration. Performance evaluation via simulation experiments reveals that the ATL protocol suite addresses the challenges posed by the NGWI and significantly improves the performance for reliable data and multimedia transport in NGWI.

***Index Terms*—Adaptive congestion control, adaptive transport layer (ATL), multimedia transport, next-generation wireless Internet (NGWI), reliable data transport.**

I. Introduction

THE SPECTACULAR evolution of the wireless network technologies has yielded several generations of wireless systems referred by nG, $n = 1, 2, 3, 4$. Currently, proposed third-generation (3G) wireless systems such as the Universal Mobile Telecommunication System (UMTS) and the International Mobile Telecommunication System 2000 (IMT-2000) can provide up to 2 Mb/s data rates along with the multimedia support [23]. These developments in the wireless networking technologies and the drastic increase in their usage have further dictated the integration of wired and wireless domains. Consequently, the next-generation wireless Internet (NGWI), which is currently in the design stage and also referred to as 4G wireless networks, will emerge as the convergence of next-generation wireless systems and Internet with the objective of *anywhere*, *anytime* seamless service to the mobile users.

A typical proposed NGWI architecture is illustrated in Fig. 1. In this architecture, an NG wireless terminal can roam between the diverse set of wireless architectures, i.e., wireless local area network (WLAN), 3G cellular, and satellite network, as shown in Fig. 1. Along with this architectural diversity, NGWI is also expected to provide the mobile users with a wide range of services from high rate data traffic to real-time multimedia traffic with a certain quality-of-service (QoS) level. These objectives pose several challenges for the realization of the NGWI as follows.

- **Heterogeneous Wireless Architectures**: The wireless systems that will be part of the NGWI have different characteristics which are summarized as follows.
  - — *Access Delay*: While the one-way access delay in the wireless link may not be significant in WLAN, a typical round-trip time (RTT) varies between few hundred milliseconds and 1 s in 3G links due to the extensive physical layer processing, e.g., forward error correction (FEC), interleaving and transmission delays [13]. The access delay is much higher in satellite links, which have high propagation delay up to 270 ms [1].
  - — *Link Errors*: The conventional transmission control protocols (TCPs) presume the existence of an underlying reliable physical link. Thus, they invoke congestion control in case of wireless link related packet losses. This leads to significant performance degradation, whose severity is proportional to the wireless link conditions, i.e., bit-error rate and delay. The packet loss rates vary from very low levels in near-wireless environments such as WLAN, 3G picocells to higher than 1% in macrocellular environments and satellite networks [27]. The mobility related packet losses, i.e., blackouts due to handoff or signal loss, amplify the degradation.
  - — *Mobility Pattern*: The mobility rate may increase the number of blackouts due to handoff and, hence, decrease the transport efficiency. During a connection

Manuscript received December 3, 2002; revised June 30, 2003. This work was supported in part by the National Science Foundation under Grant ANI-0117840.

The authors are with the Broadband and Wireless Networking Laboratory, School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332 USA (e-mail: akan@ece.gatech.edu; ian@ece.gatech.edu).

Digital Object Identifier 10.1109/JSAC.2004.826919

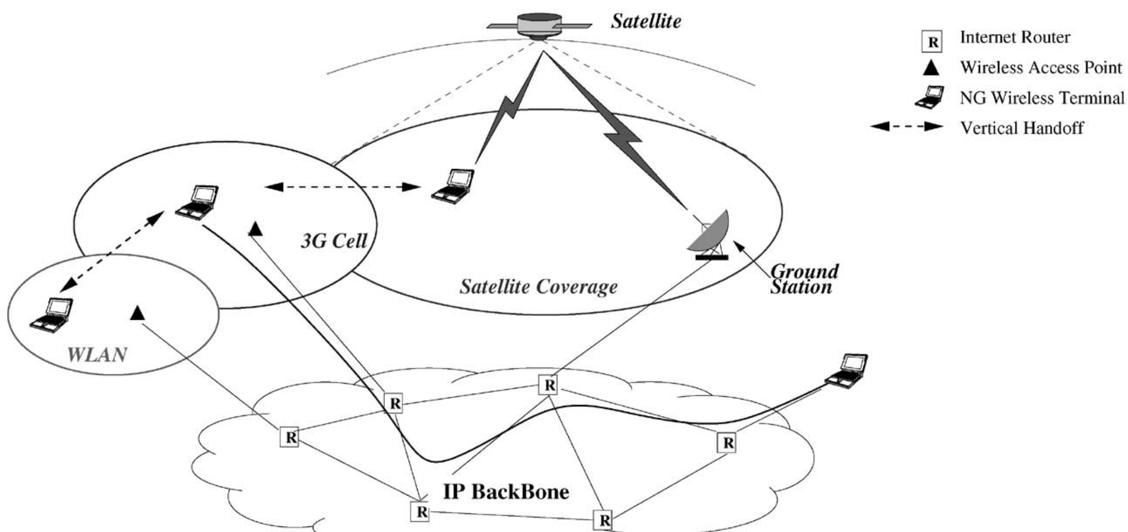

Fig. 1. Typical proposed NGWI architecture including WLAN, 3G, and satellite networks converged with the next-generation Internet backbone.

period, almost no handoff is experienced in the global coverage, whereas frequent handoffs may take place in picocellular environments.

- **Heterogeneous Service Demands**: The services that will be provided by the NGWI vary from high rate reliable data to real-time multimedia such as live video streaming. While the former requires 100% reliable transport, the latter, instead, needs timely delivery and smooth rate variation. For multimedia traffic delivery, user datagram protocol (UDP)-based streaming without any transmission rate control may also lead to unfairness to the TCP sources and further result in a congestion collapse [8].

These challenges need to be addressed to realize the NGWI objective. The architectural heterogeneities described above need to be handled dynamically, while the mobile user migrates during connection period. The diversity in the demanded services necessitates both the reliable data transport functionality and the rate control protocol for time-constrained multimedia delivery. All these heterogeneities coupled with the processing and power limitations of wireless terminals call for a unified, efficient, and seamless *adaptive transport layer*. In order to address this need, a new adaptive transport layer (ATL) for the NGWI is presented in this paper.

The ATL is a unified adaptive transport layer that *dynamically adjusts its protocol configurations to adapt to heterogeneous wireless environments and supports reliable data and multimedia delivery*. Some of its salient features are as follows.

1) *Adaptive Congestion Control*: ATL incorporates a new adaptive congestion control for its new reliable data transport protocol (TCP-ATL) and its new rate control protocol for multimedia traffic (RCP-ATL). The additive-increase multiplicative-decrease (AIMD) control parameters of both protocols are dynamically adjusted according to the current wireless link conditions, i.e., the packet loss rate and the access delay. This enables ATL to achieve high-throughput performance despite the architectural diversity of the NGWI. The adaptive congestion control scheme described in Section III also implicitly addresses the effects of blackouts.
2) *Multimedia Support*: A new UDP-based rate control protocol (RCP-ATL) for multimedia traffic in the NGWI is also embodied within the ATL. RCP-ATL performs rate control scheme by using the adaptive congestion control to achieve high goodput performance in the heterogeneous wireless architectures. RCP-ATL also cooperates with an adaptive encoder to achieve homogeneous and smooth quality variation as explained in Section IV-B
3) *Fairness*: The methods for dynamic adjustment of the AIMD parameters of adaptive congestion control for both TCP-ATL and RCP-ATL are developed in Section III by taking the fairness into consideration. Therefore, ATL protocols inherently preserve fairness to the wired TCP sources sharing the same bottleneck.
4) *Low Complexity and Backward Compatibility*: The ATL protocols are developed based on the current TCP and UDP protocols with additional functionalities tailored to address the unique requirements of the NGWI. Therefore, ATL can easily communicate with current transport layer implementations, which follow conventional TCP/IP semantics. This approach also matches with a unified, low-complexity adaptive transport layer requirement for the processing and memory constrained mobile terminals.

The remainder of the paper is organized as follows. A review of related work in transport layer protocols for wireless systems and their inadequacies in addressing the diversity in the NGWI are presented in Section II. The effects of the architectural heterogeneity on the performance are investigated via a case study and the adaptive congestion control methods for both TCP-ATL and RCP-ATL are then developed in Section III. The overview of ATL protocols along with their detailed operations are explained in Section IV. The performance evaluation of TCP-ATL and RCP-ATL and the simulation results are presented in Section V. Finally, the paper is concluded in Section VI.

II. Related Work

Despite the extensive research in transport protocols to address the challenges in wireless domain [1]–[4], [6], [10], [25], most of them address these challenges only for a specific wireless environment and, hence, do not solve the problems pertaining to the NGWI. For example, independent solutions have been proposed for microcell [3] and macrocell [25] environments, and for satellite networks [1], [2], [10]. The Snoop protocol [3] proposed primarily for WLANs has been shown in [25] to be inefficient in 3G environments due to its assumption of insignificant wireless link delays compared with the end-to-end delays. On the other hand, WTCP proposed in [25] for wireless wide area networks (WWAN) relies on the interpacket separation as a congestion metric. Hence, WTCP is not applicable to WLAN and picocells, since its congestion metric is not reliable in the high-bandwidth low-delay environments. In [6], TCP-Westwood is proposed to improve the TCP performance in the hybrid wired/wireless networks. TCP-Westwood uses acknowledgment (ACK) reception rate to estimate the available bandwidth, which is then used to calculate the congestion window and slow-start threshold. However, for the links with high delay, TCP-Westwood performance degrades due to the decrease in the estimation accuracy because of late arriving feedback [6]. Furthermore, the abrupt changes in wireless link delay due to vertical handoffs between different wireless systems of the NGWI may also decrease the accuracy of its bandwidth estimation method.

Similarly, the solutions for satellite networks also do not stand as a unified solution to address the heterogeneities in the NGWI. TCP-Peach [1] and its enhanced version TCP-Peach+ [2] significantly improve the throughput performance in the satellite links. However, their congestion control algorithms, which are tailored to match requirements of the links with high-bandwidth-delay products, may not be applicable to WLAN and picocell environments.

For the NGWI architecture in which the mobile user migrates between the different environments as illustrated in Fig. 1, none of these protocols will suffice. In [11], an end-to-end transport layer approach called pTCP is proposed to perform bandwidth aggregation for mobile terminals that are connected to multiple wireless networks. pTCP opens and maintains one TCP-virtual connection for every interface to perform data striping. Although pTCP does not propose specific congestion control schemes to be used for each of the interfaces, it has been pointed out in [11] that incorporating interface-specific transport protocols to address the architectural heterogeneity could be a conceivable solution until a unified transport layer framework is developed. However, including each of the existing transport protocols tailored for a specific architecture in a single transport layer is not a practical solution due to the processing and memory constraints of wireless terminals. Therefore, a unified and adaptive transport protocol that adapts itself to the different wireless environments has yet to be developed to address the architectural heterogeneities posed by the NGWI objective.

On the other hand, there exists a necessity of a rate control mechanism for multimedia flows to avoid unfairness to TCP sources and further congestion collapse [8]. Although there exists a significant amount of research in this context [5], [7], [15], [21], [22], most of them are proposed for wired networks. These solutions follow the conservative rate halving behavior of TCP, which leads to unnecessary rate throttle and, hence, severe performance degradation in the wireless environments. In [29], an end-to-end wireless multimedia streaming TCP-Friendly protocol (WMSTFP) is proposed for multimedia over wireless Internet. In WMSTFP, the sender adjusts the sending rate by calculating the available bandwidth estimation based on the feedback obtained from the receiver. Although the WMSTFP improves the network utilization and maintain fairness to the TCP sources, it is not developed to address all of the challenges in the heterogeneous architectures of the NGWI. For example, in the satellite networks, the RTT-based rate estimation technique used in WMSTFP may not achieve high utilization of the link resources due to very high propagation delay. In [27], rate control scheme (RCS) has been proposed for real-time traffic in the networks with high-bandwidth-delay products and high bit-error rates. RCS significantly improves the throughput performance, while maintaining fairness. However, its *dummy packet* based congestion control algorithm is specifically tailored to match the requirements of the links with high propagation delay and may be inefficient for the wireless environments with low access delay such as WLAN and picocells.

Despite the significant amount of research in transport protocols for wireless networks and rate control schemes for multimedia traffic flows, there is no single solution to address all of these heterogeneities pertaining to the NGWI. Hence, there exists a need for a unified transport layer to handle both architectural diversity and heterogeneous application requirements of the NGWI.

III. Adaptive Congestion Control

In this section, the variation in the performance degradation severity due to the architectural heterogeneity in the NGWI is first investigated via a case study. Then, the adaptive congestion control for reliable data transport and multimedia traffic delivery are developed.

A. Architectural Heterogeneity: Case Study

The performance degradation experienced by current TCP protocols is primarily affected by the wireless link conditions. The increase in the probability p_w of packet loss in the wireless link also increases the severity of the performance degradation. Similarly, an increase in the RTT of the end-to-end path due to an additional wireless link delay d_w further degrades the network utilization [1]. Moreover, the high RTT also prevents terminal from utilizing its fair bandwidth share because of the well-known unfairness to the TCP flows with higher RTT [14]. Therefore, the severity of the performance degradation changes according to the varying wireless link conditions.

In order to study the effects of the architectural heterogeneity in the NGWI on the TCP performance, we performed simulation experiments using $ns-2$ [19]. Here, three wireless and one wired TCP-Newreno sources are connected to four receivers via a 2.5-Mb/s bottleneck link. The wireless TCP sources are connected to the bottleneck via links with the capacity of 1 Mb/s, the packet loss probabilities of $p_w = 10^{-4}, 10^{-3}, 10^{-2}$, and the

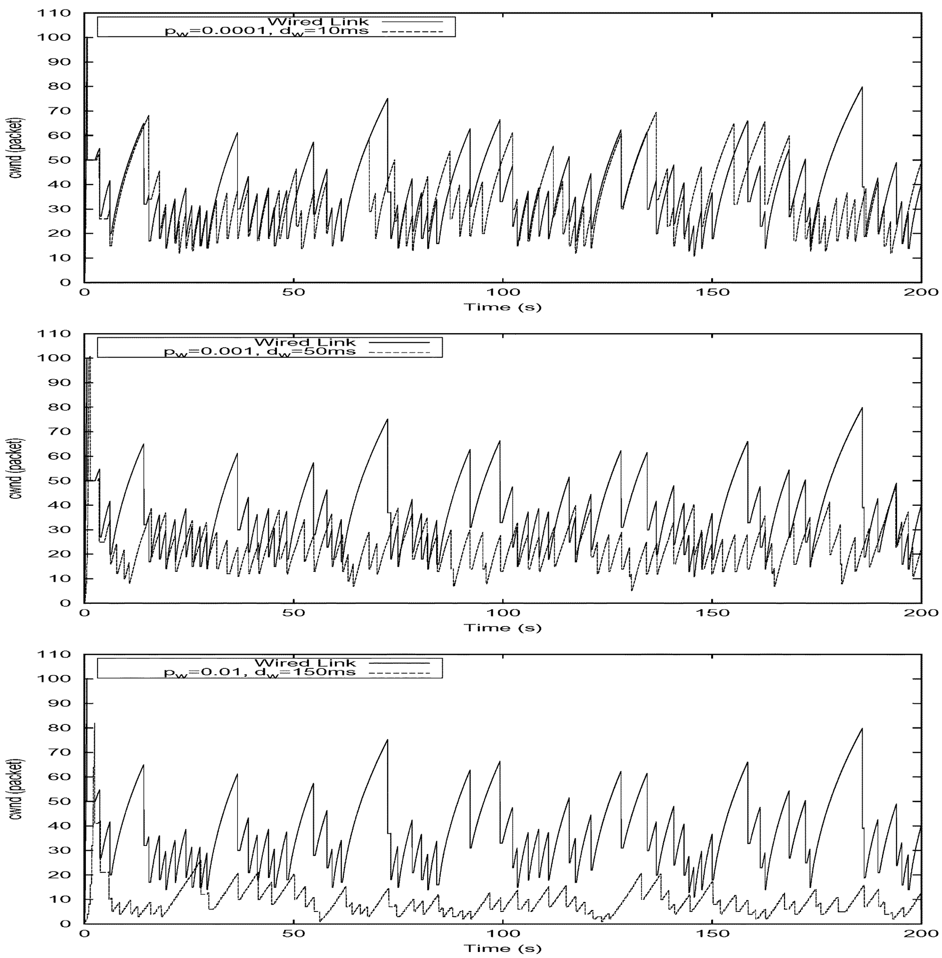

Fig. 2. Congestion window change with time for different wireless link conditions.

TABLE I
SIMULATION PARAMETERS AND RESULTS

Wireless System	d_w (ms)	p_w	*B/W (Mb/s)*	*Tput (Mb/s)*	*Gput (Mb/s)*
Wired	10	0	1	0.7286	0.86555
WLAN	10	10^{-4}	1	0.7271	0.84686
3G Cellular	50	10^{-3}	1	0.4899	0.57977
Satellite	150	10^{-2}	1	0.1549	0.12417

one-way link delays of $d_w = 10$, 50, and 150 ms, respectively. All TCP sources are attached with a file transfer protocol (FTP) application. Note that three (p_w, d_w) pairs used in the simulation experiments represent the typical wireless link conditions in WLAN, 3G cellular, and satellite networks.

The congestion window $(cwnd)$ change of the TCP sources are shown in Fig. 2. The wired TCP source halves its congestion window only in case of packet loss due to congestion. However, the wireless TCP sources perform more rate throttles, which is proportional to the p_w as shown in Fig. 2. Furthermore, the recovery from $cwnd$ halving takes time proportional to the RTT, which increases with the wireless link delay d_w. While the performance of the wireless source connected to WLAN is very close to the wired TCP source, the TCP source representing a satellite connection almost cannot utilize its entire link, as shown in Fig. 2. The parameters and the results of the simulation experiments are summarized in Table I.

We observe in Fig. 2 and Table I that the achievable throughput performance significantly varies with varying wireless environment conditions. The same experiments are also conducted with an AIMD rate control protocol [22] for multimedia traffic and similar patterns have been observed, and their goodput results are also given in Table I. These results motivate us to develop a unified and adaptive transport layer protocols that maintain the maximum achievable throughput in all wireless architectures of the NGWI.

B. Adaptive Congestion Control for Reliable Data Transport

It is shown in Section III-A that the TCP performance significantly varies with wireless link conditions. Therefore, a new adaptive congestion control is necessary to maintain the highest throughput performance possible in all wireless architectures. On the other hand, this objective is upper bounded with the throughput of the wired TCP source sharing the same bottleneck in order to maintain the fairness. For example, the $cwnd$ curve achieved by the wired TCP in Fig. 2 represents the target throughput to achieve for the other three wireless resources. Therefore, the objective for adaptive congestion control can be restated as the minimization of the difference between the areas under the $cwnd$ curves achieved by the wired TCP and the wireless TCP sources.

As observed in Fig. 2, the TCP throughput directly depends on the link conditions. Let $\mathcal{T}_{\alpha,\beta}(p, R, T_0, b)$ be the throughput achieved by a TCP source, where (α, β) are its additive-increase and multiplicative-decrease parameters, p is the packet loss probability, R is the end-to-end RTT, T_0 is the initial retransmission timeout (RTO), and b is the number of data packets acknowledged with a single ACK. Hence, the wired TCP throughput in Fig. 2 can be expressed by $\mathcal{T}_{1,1/2}(p_c, R_c, T_{0c}, b)$, where p_c is the probability of packet loss experienced due to congestion, and R_c is the end-to-end RTT of the connection without any additional wireless link delay. Therefore, $\mathcal{T}_{1,1/2}(p_c, R_c, T_{0c}, b)$ is the upper bound for the target throughput due to fairness consideration. Hence, the objective of the adaptive congestion control method is to achieve $\mathcal{T}_{1,1/2}(p_c, R_c, T_{0c}, b)$, regardless of the underlying wireless architecture and the link conditions.

In the simulation experiments presented in Section III-A, the wireless TCP sources achieved lower throughput than the wired TCP source due to the increased packet loss rate and RTT because of the additional packet losses and the link delay experienced in the wireless link. This situation can be generalized for $p > p_c$ and $R > R_c$ as

$$\mathcal{T}_{1,\frac{1}{2}}(p, R, T_0, b) < \mathcal{T}_{1,\frac{1}{2}}(p_c, R_c, T_{0c}, b) \tag{1}$$

where p and R are the total probability of packet loss due to both wireless link and congestion and the end-to-end RTT, respectively. Here, p and R can be expressed by

$$p = 1 - (1 - p_w)(1 - p_c) \tag{2}$$

$$R = R_c + 2d_w \tag{3}$$

where p_w is the probability of wireless link related packet loss, p_c is the probability of packet loss due to congestion, R_c is the end-to-end RTT of the connection without any additional wireless link delay, and d_w is the one-way wireless link delay as in Table I.

As the mobile terminal migrates between different wireless environments of the NGWI, p_w and d_w and, hence, p and R in (2) and (3), vary continuously. Such variation in the link conditions leads the mobile terminal to experience different throughput degradation levels, as observed in Fig. 2.

To maintain the highest throughput performance in all of the different wireless architectures characterized by different access delays and link error rates and achieve fairness to the wired TCP sources sharing the same bottleneck, ATL incorporates a new adaptive congestion control that dynamically adjusts its AIMD parameters (α, β) according to the current wireless link conditions, i.e., p_w and d_w. Therefore, the objective of the new adaptive congestion control is to obtain (α, β) pair such that

$$\mathcal{T}_{\alpha,\beta}(p, R, T_0, b) \approx \mathcal{T}_{1,\frac{1}{2}}(p_c, R_c, T_{0c}, b) \tag{4}$$

is achieved in all of the different architectures of the NGWI; i.e., $\forall p_w$ and $\forall d_w$ for $p > p_c$ and $R > R_c$, where p and R are expressed in (2) and (3) as a function of wireless link condition parameters p_w and d_w.

Let $\widehat{T}$ be the upper bound for the target throughput to be achieved by the wireless TCP sources, i.e., $\widehat{T} = \mathcal{T}_{1,1/2}(p_c, R_c, T_{0c}, b)$. Recall that this upper bound is imposed by the fairness consideration, and hence, $\widehat{T}$ is the throughput achieved by the wired TCP source experiencing p_c and R_c as the packet loss probability due to congestion and the end-to-end RTT (with no additional wireless link delay), respectively. Given that p, R, p_w, and d_w are known, it follows from (2) and (3) that p_c and R_c can be calculated as

$$p_c = \frac{p - p_w}{1 - p_w} \tag{5}$$

$$R_c = R - 2d_w. \tag{6}$$

The throughput of the TCP connection $\mathcal{T}_{\alpha,\beta}(p, R, T_0, b)$ in (4) can be expressed by [28]

$$\mathcal{T}_{\alpha,\beta}(p, R, T_0, b) = \frac{1}{R\sqrt{\frac{2b(1-\beta)p}{\alpha(1+\beta)}} + T_0\left(3\sqrt{\frac{(1-\beta^2)bp}{2\alpha}}\right)p(1+32p^2)} \tag{7}$$

where (α, β) are the additive-increase and multiplicative-decrease parameters, p is the packet loss probability, R is the end-to-end RTT, T_0 is the initial retransmission timeout (RTO), and b is the number of data packets acknowledged with a single ACK. Therefore, $\widehat{T}$ can be calculated by feeding (7) with p_c, R_c, T_{0c}, and b. Here, each data packet can be assumed to be acknowledged by an ACK, i.e., $b = 1$, and T_{0c} can be approximately set to $4R_c$ to provide fairness with TCP [8]. Therefore, by substituting the obtained $\widehat{T}$ and (7) into (4), we get

$$R\sqrt{\frac{2b(1-\beta)p}{\alpha(1+\beta)}} + T_0\left(3\sqrt{\frac{(1-\beta^2)bp}{2\alpha}}\right)p(1+32p^2) = \frac{1}{\widehat{T}}. \tag{8}$$

Consequently, the additive-increase parameter α that achieves the objective stated in (4) can be calculated from (8) as

$$\alpha = \frac{bp(1-\beta)}{2(1+\beta)}\left[\widehat{T}\left(2R + 3T_0p(1+32p^2)(1+\beta)\right)\right]^2. \tag{9}$$

Therefore, once p_w, d_w, p, and R are known, $\widehat{T}$ can be calculated by (7); then, (9) can be used to set the additive-increase parameter α for $\forall\beta \in [0.5, 1)$. The selection of the multiplicative-decrease factor β from [0.5,1) is discussed in Sections IV-A and V-A1. By this way, the throughput degradation can be avoided and highest throughput, which is upper bounded by the fairness constraint, is achieved via dynamically adjustment of AIMD parameters according to the wireless link conditions. The detailed operation of the adaptive congestion control for reliable data transport is presented in Section IV-A.

C. *Adaptive Rate Control for Multimedia Traffic*

Similarly, a new adaptive rate control for multimedia flows is necessary to maintain the highest performance possible in all of the different wireless architectures, while preserving fairness to the TCP sources. As in Section III-B, the highest throughput is again upper bounded by the throughput of the wired TCP source sharing the same bottleneck in order to maintain fairness.

Assume a general *rate-based* AIMD scheme which increases the transmission rate S additively with α at each RTT, i.e., $S = S + \alpha$, and reduces it multiplicatively by β if a packet loss is detected, i.e., $S = S \cdot \beta$. The steady-state throughput of such *rate-based* AIMD scheme is derived (see the Appendix) and expressed by

$$\mathcal{T}^{r}_{\alpha,\beta}(p,R)=\frac{\alpha}{4(1-\beta)}\left[1+\beta+\sqrt{(3-\beta)^2+\frac{8(1-\beta^2)}{\alpha Rp}}\right] \tag{10}$$

where p is the packet loss probability, R is the RTT, and α and β are the additive-increase and the multiplicative-decrease factors, respectively.

It is observed from Table I and (10) that the throughput of the rate-based AIMD scheme depends on the link conditions, i.e., p and R, and the values of α and β. In order to maintain the highest throughput performance and achieve fairness to the wired TCP sources, ATL adapts the AIMD parameters of its rate control protocol to the link conditions, i.e., p_w and d_w. As in Section III-B, the objective of the new adaptive rate control is to obtain (α,β) pair such that

$$\mathcal{T}^{r}_{\alpha,\beta}(p,R)\approx\mathcal{T}_{1,\frac{1}{2}}(p_c,R_c,T_{0c},b) \tag{11}$$

is achieved in all of the different architectures of the NGWI, i.e., $\forall p_w$ and $\forall d_w$ for $p>p_c$ and $R>R_c$, where p and R are expressed in (2) and (3) as a function of wireless link condition parameters p_w and d_w.

Let $\widehat{T}$ be the upper bound for the target throughput to be achieved by the wireless multimedia traffic sources as in Section III-B, i.e., $\widehat{T}=\mathcal{T}_{1,1/2}(p_c,R_c,T_{0c},b)$. Thus, by substituting the $\widehat{T}$ and (10) into (11), we obtain

$$\frac{\alpha}{4(1-\beta)}\left[1+\beta+\sqrt{(3-\beta)^2+\frac{8(1-\beta^2)}{\alpha Rp}}\right]=\widehat{T}. \tag{12}$$

Hence, the additive-increase parameter α of the AIMD-based rate control protocol to achieve a target throughput of $\widehat{T}$ can be obtained as

$$\alpha=\frac{(1+\beta)}{2}\left(\widehat{T}+\frac{1}{R\cdot p}\right)\left[\sqrt{1+\frac{8\widehat{T}^2(1-\beta)}{\left(\widehat{T}+\frac{1}{R\cdot p}\right)^2(1+\beta)^2}}-1\right]. \tag{13}$$

Therefore, given that p_w, d_w, p, and R are known, then (13) can be used to set α for $\forall\beta\in[0.5,1)$ according to the wireless link conditions. The selection of the multiplicative-decrease factor β from [0.5,1) is discussed in Sections IV-A and V-A1. By this way, an adaptive rate control scheme can achieve the highest possible throughput, while maintaining fairness. The detailed operation of the adaptive rate control scheme for multimedia traffic is explained in Section IV-B.

IV. Adaptive Transport Layer (ATL) for NGWI

ATL is a unified adaptive transport layer that incorporates a new reliable data transport protocol (TCP-ATL) and a new rate control protocol (RCP-ATL) for multimedia traffic. The protocol structure of ATL is illustrated in Fig. 3. According to the service type requested by the application layer or the incoming request from the receiver side, the appropriate protocol algorithm is initiated. The application layer interfacing and the connection establishment procedures are left out of the scope of this paper. In this section, the operations of the TCP-ATL and the RCP-ATL are presented in detail.

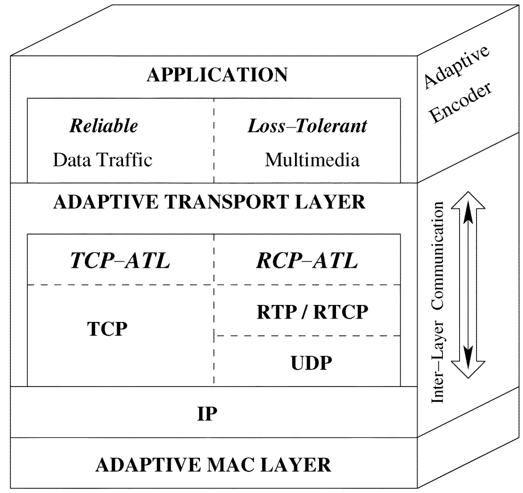

Fig. 3. Typical protocol stack including ATL.

A. TCP-ATL: Adaptive Reliable Data Transport Protocol

ATL incorporates a new adaptive reliable transport protocol (TCP-ATL) to address the architectural heterogeneity of the NGWI. TCP-ATL inherits well-established TCP functionalities of the legacy TCP protocols [26] and utilizes the adaptive congestion control scheme developed in Section III-B. TCP-ATL can be also implemented based on any conventional AIMD TCP protocol to achieve the adaptive congestion control objective. Furthermore, the selective acknowledgment (SACK) options [18] are shown to be very effective in recovering multiple packet losses in one TCP window. This is much more important on the links with high bandwidth-delay-product such as satellite links. Therefore, TCP-ATL uses the SACK options [18] to guarantee reliability.

In Section III-B, it is shown that the AIMD parameters of TCP can be adjusted according to the link conditions by using (9) given that p, R, p_w, and d_w are known. To utilize (9) for this goal, the TCP-ATL source continuously measures the packet loss rate p within a sliding time window of τ. The wireless link delay d_w and probability of wireless related packet loss p_w are assumed to be obtained from the underlying *adaptive medium access control (MAC) layer*, which is also essential to address the architectural heterogeneities in the NGWI. In practice, the wireless MAC protocols have the information about the packet loss rate and the delay experienced in the wireless access link, i.e., p_w and d_w. The access delay is immediately available at the MAC layer, since many MAC protocols use it to calculate their timeout values, e.g., clear-to-send (CTS) timeout value calculation in IEEE 802.11 [12]. Furthermore, wireless MAC protocols perform several attempts to get the packets it receives from upper layer to the access point. If the MAC layer cannot successfully get the data packet across, it informs the upper layer about this result. Note that this failure may occur because of several wireless link conditions, such as access contention, link errors, link outage, etc. Hence, the packet loss rate due to wireless link information p_w can be obtained from the MAC layer information via simple cross-layer interaction. The details of such communication between different networking layers are software implementation issues and left beyond the scope of this paper. Moreover, note that p_w is also measured within a sliding time window of τ and includes all wireless link related packet losses encountered by the underlying MAC layer due to access failure, bit-errors, fading, and signal loss due to handoff or blackout.

Once p, R, p_w, and d_w are known, p_c and R_c can be obtained by using (5) and (6). Hence, $\widehat{T} = \mathcal{T}_{1,1/2}(p_c, R_c, T_{0c}, b)$ can be calculated using (7) and plugged into (9) to determine α for a given β for $\forall\beta \in [0.5, 1)$. While the selection of β affects the instantaneous TCP-friendliness [9] of the TCP-ATL behavior, it does not affect the overall fairness. This is because any (α, β) pair determined by (9) inherently complies with fairness requirement, since (9) itself is obtained by taking that requirement into consideration. Therefore, β can be safely selected from [0.5,1). On the other hand, the higher the β is selected, the higher performance can be observed in the links with high-bandwidth-delay products. This is mainly because the throughput degradation due to frequent rate throttles is amplified by the high-bandwidth-delay product of the link. Hence, selecting higher β values from [0.5,1) helps compensate for the adverse effects of high propagation delay on the performance. During the simulation experiments presented in Section V, β values of 0.75, 0.80, and 0.85 are found to give the best performance, for $d_w = 10$, 50, and 150 ms, respectively. Therefore, β can be set to be 0.75, 0.80, and 0.85, for WLAN ($0 \text{ ms} < d_w < 50 \text{ ms}$), 3G cellular ($50 \text{ ms} \leq d_w < 150 \text{ ms}$), and satellite networks ($d_w \geq 150$) ms, respectively.

In the beginning of the connection, there is no information in terms of p and p_w. Therefore, the connection starts with the default AIMD parameters, i.e., $(\alpha, \beta) = (1, 0.5)$. Note that in the absence of a wireless link, i.e., $p_w = 0$ and $d_w = 0$, the connection simply uses normal TCP which is also desirable to maintain fairness. As packet losses are being experienced due to both congestion and wireless link, (α, β) is adjusted to adapt the protocol configuration to the varying wireless link conditions. Furthermore, at each vertical handoff, it is essential that ATL updates its AIMD parameters to address the possible abrupt variation in the wireless link conditions. It is assumed that the vertical handoff events are informed by underlying MAC layer over interlayer communication plane shown in Fig. 3.

On the other hand, the blackout situations are also implicitly addressed by the adaptive congestion control algorithm of TCP-ATL. If a blackout occurs due to mobility or fading by the environmental obscurations, this immediately increases p_w. The TCP-ATL source then calculates new α using (9) to adapt the congestion control configuration to the new link conditions. It follows from (9) that an increase in p_w also increases α. Therefore, TCP-ATL can quickly recover from the throughput degradation due to a blackout by increasing its congestion window more aggressively. Once the signal is back, p_w immediately drops which also reduces α. Therefore, TCP-ATL can dynamically return to its configuration just before the blackout.

The default TCP-ATL operation explained above directly applies to the scenarios, where the source is wireless and the receiver is wired terminals. The TCP-ATL operation may require slight modifications in terms of calculations of p_c and R_c according to the different source/receiver combinations. The operation of the protocol for each source/receiver combination is summarized as follows.

1) *Wireless Source/Wired Receiver*: In this case, the default ATL operation explained in this section does not require any modifications.
2) *Wired Source/Wireless Receiver*: In this case, the wireless link parameters, i.e., p_w and d_w, are supplied to the source. Once the wireless terminal acquires p_w and d_w for the first time, it forwards them to the source within data ACK packets to avoid any additional overhead. p_w and d_w are then forwarded to the source if they vary. Similarly, the vertical handoff event is also signaled to the source using data ACK packets in this scenario.
3) *Wireless Source/Wireless Receiver*: In this case, the end-to-end path includes two wireless links. Let p_w^s, d_w^s and p_w^r, d_w^r be the packet loss probability and the delay of the wireless link at the source and the receiver side of the end-to-end path, respectively. As the wireless receiver first obtains p_w^r and d_w^r, it forwards them to the source as explained above. The wireless ATL source, which receives this information, then infers that the receiver is also wireless terminal equipped with the ATL. In this case, the total packet loss probability p and the RTT R, which are previously expressed by (2) and (3), can be calculated as

$$p = 1 - (1 - p_w^s)(1 - p_c)(1 - p_w^r) \tag{14}$$
$$R = R_c + 2(d_w^s + d_w^r). \tag{15}$$

It follows from (14) and (15) that the wireless source, which measures p and R as explained before, can calculate p_c and R_c as

$$p_c = 1 - \frac{1 - p}{(1 - p_w^s)(1 - p_w^r)} \tag{16}$$
$$R_c = R - 2(d_w^s + d_w^r). \tag{17}$$

Once p_c and R_c are found using (16) and (17), the original TCP-ATL operation is then resumed.
4) *Wired Source/Wired Receiver*: In this case, no wireless link involves in the entire connection path, that is, $p_w = 0$ and $d_w = 0$. Therefore, normal TCP protocol operation inherited by TCP-ATL resumes the connection. Note that ATL is also backward compatible as it can still establish TCP connection even if any end-party is not equipped with the ATL.

The operation of the adaptive congestion control scheme used by TCP-ATL protocol is summarized in the pseudoalgorithm shown in Fig. 4.

B. RCP-ATL: Adaptive Rate Control Protocol for Multimedia Traffic

ATL incorporates a new adaptive rate control protocol (RCP-ATL) to address the architectural heterogeneity of the NGWI. RCP-ATL is an end-to-end rate control protocol that uses the adaptive rate control scheme developed in Section III-C

```
Configure_AIMD()
  /* Case 1 */
  if (Wireless Source / Wired Receiver)
    Obtain p_w and d_w from MAC layer;
    p_c = (p - p_w)/(1 - p_w);
    R_c = R - 2d_w;
  end;
  /* Case 2 */
  if (Wired Source / Wireless Receiver)
    Obtain p_w and d_w from RECEIVER;
    p_c = (p - p_w)/(1 - p_w);
    R_c = R - 2d_w;
  end;
  /* Case 3 */
  if (Wireless Source / Receiver)
    Obtain p_w^s and d_w^s from MAC layer;
    Obtain p_w^r and d_w^r from RECEIVER;
    d_w = d_w^s + d_w^r;
    p_c = 1 - (1 - p)/((1 - p_w^s)(1 - p_w^r));
    R_c = R - 2(d_w^s + d_w^r);
  end;
  Calculate T^ = T_{1,1/2}(p_c, R_c, T_0c, b) with (7);
  if (0 ms < d_w < 50 ms)
    Set β = 0.75;
  if (50 ms ≤ d_w < 150 ms)
    Set β = 0.80;
  if (d_w ≥ 150 ms)
    Set β = 0.85;
  if (TCP_ATL)
    Set α with (9);
  else if (RCP_ATL)
    Set α with (13);
  end;
end;
```

Fig. 4. Pseudoalgorithm for the operation of the adaptive congestion and rate control schemes used by TCP-ATL and RCP-ATL protocols.

in order to produce TCP-friendly traffic flows, while maintaining high-throughput performance in the NGWI.

RCP-ATL can run on top of RTP/RTCP [24] and UDP as shown in Fig. 3. RCP-ATL is not an ARQ protocol and the source does not perform retransmission due to tighter time constraints of the multimedia flows. However, the receiver sends back an ACK for any received packet. If the RTP/RTCP is implemented, then these ACKs can be the receiver reports (RR) that includes the reception quality statistics such as the number of packets received, RTP timestamp, fraction lost, and cumulative number of packets lost. Hence, RCP-ATL can obtain p and R from the RTP receiver reports. If RTP/RTCP is not used for any reason, then data ACKs would be sufficient to acquire p and R.

RCP-ATL follows the AIMD rate control behavior. It multiplicatively decreases its data rate (S) in case of a packet loss, i.e., $S = S \cdot \beta$. Otherwise, it additively increases the data rate with α at each RTT, i.e., $S = S + \alpha$. Furthermore, RCP-ATL can be implemented based on any existing AIMD rate control protocol to achieve the adaptive congestion control objective.

As in Section IV-A, p_w and d_w are assumed to be obtained from the underlying *adaptive MAC layer*. Once p, R, p_w, and d_w are known, p_c and R_c can be obtained by using (5) and (6). Hence, $\widehat{T} = \mathcal{T}_{1,1/2}(p_c, R_c, T_{0c}, b)$ can be calculated using (7) and plugged into (13) to determine α for a given β for $\forall\beta \in [0.5, 1)$. The discussions on the selection of β given in Section IV-A also apply here. In this case, the selection of β also affects the smoothness of the rate variation, which is essential to maintain a certain quality level in multimedia streaming. Although any (α, β) pair determined by (13) inherently complies with fairness requirement, higher β is preferable for smoothness requirement. The investigation of the effects of the β selection on the smoothness of the rate variation is left for future study.

On the other hand, RCP-ATL rate control protocol can be in cooperation with an adaptive media encoding process. The adaptive encoder can be provided with the available bandwidth $S(t)$ estimated by the RCP-ATL. Hence, the real-time multimedia encoder can then adapt its encoding rate $R(t)$ according to the controlled fair share of the network resources, i.e., $R(t) \approx S(t)$. By this way, the quality of the encoded multimedia can vary homogeneously and smoothly without leading to congestion and undesirable abrupt quality variations.

As explained for TCP-ATL in Section IV-A, RCP-ATL also indirectly addresses the blackout situations by its adaptive rate control algorithm. On the other hand, the RCP-ATL may require simple modifications in terms of calculations of p_c and R_c according to the different source/receiver combinations. The explanation of the TCP-ATL operation given in Section IV-A for each source/receiver combination also applies to RCP-ATL and, hence, is not repeated here. The only difference is that p_w, d_w, and vertical handoff event information can now be forwarded to the source within RTP receiver reports rather than data ACKs. However, if RTP/RTCP is not used for any reason, p_w, d_w, and vertical handoff event information can be still sent to the source within data ACKs. The operation of the adaptive rate control scheme used by RCP-ATL protocol is also summarized in the pseudoalgorithm shown in Fig. 4.

V. PERFORMANCE EVALUATION

In order to investigate the performance of the ATL, extensive simulation experiments are conducted. The throughput, blackouts, and fairness performance of TCP-ATL in the NGWI are evaluated in Section V-A. In Section V-B, the performance of RCP-ATL in terms of throughput, fairness, and jitter is investigated.

A. TCP-ATL Performance

1) Throughput: TCP-ATL simulation experiments are performed for varying packet loss probability (p_w) between 10^{-6} and 10^{-1}. The experiments are conducted for three different one-way wireless link delay values, i.e., $d_w =$ 10, 50, and 150 ms, which represent typical wireless link delay values in WLAN, 3G cellular, and satellite networks, respectively. During these three sets of experiments, the multiplicative-decrease factor β is set to be 0.75, 0.80, and 0.85, respectively. These settings are found to give the best performance for simulation experiments. This makes perfectly sense since the throughput degradation (due to frequent rate throttles) increases with the bandwidth-delay product of the link. Therefore, the higher the β is selected, the higher performance is observed in the links with high-bandwidth-delay products. The throughput values are obtained at the sender. The simulations are run on the following network configuration. The source and the destinations are connected to each other via two routers connected

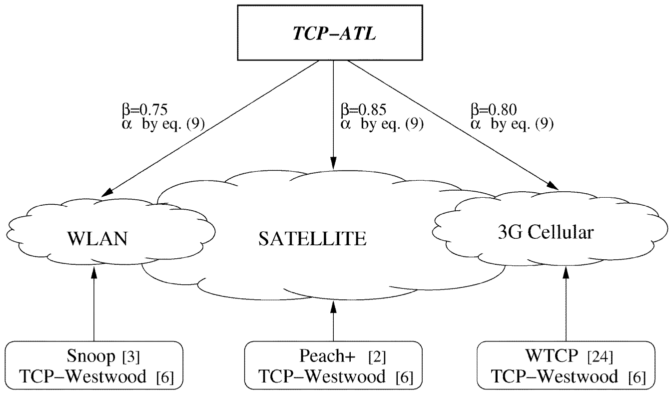

Fig. 5. Unified and adaptive TCP-ATL protocol instead of using different transport protocols that are specifically designed for different wireless architectures.

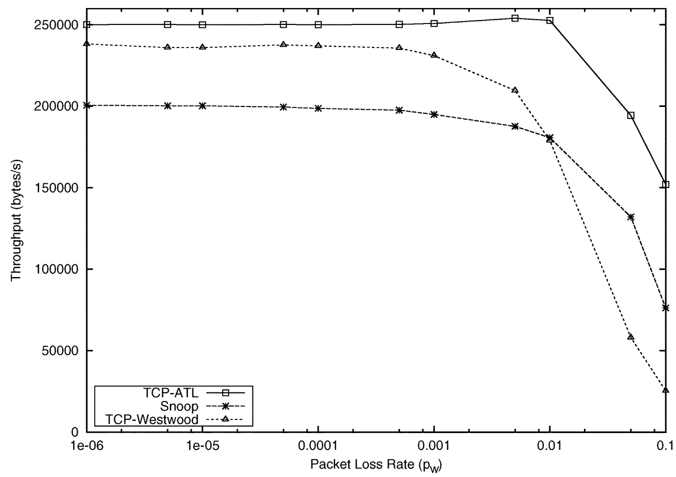

Fig. 6. Throughput comparison of TCP-ATL, Snoop, and TCP-Westwood for $d_w = 10$ ms and varying p_w.

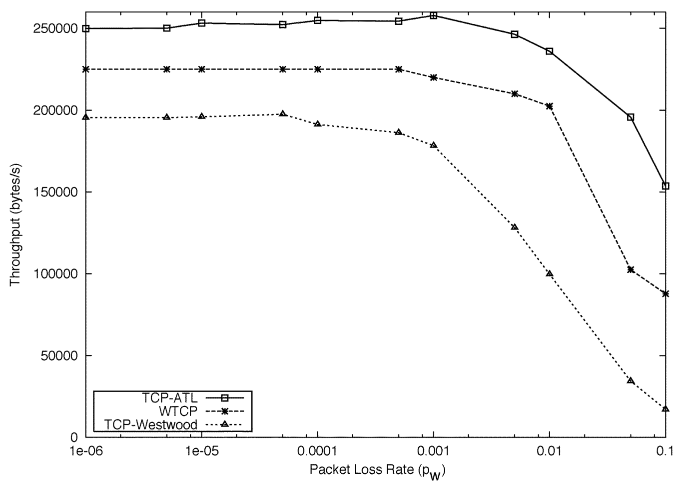

Fig. 7. Throughput comparison of TCP-ATL, WTCP, and TCP-Westwood for $d_w = 50$ ms and varying p_w.

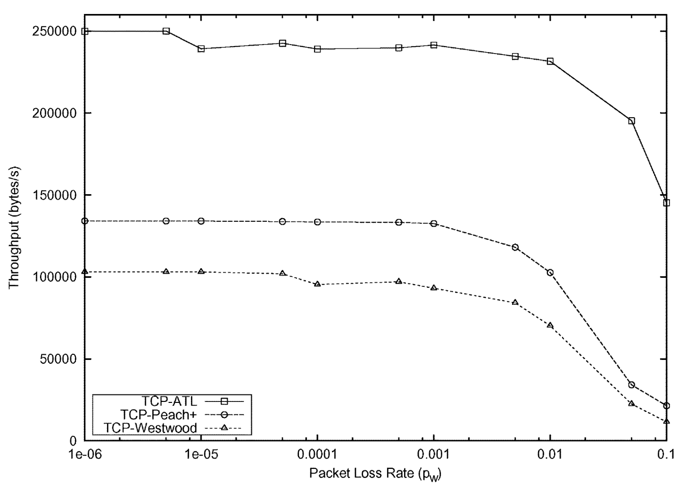

Fig. 8. Throughput comparison of TCP-ATL, TCP-Westwood, and TCP-Peach+ for $d_w = 150$ ms and varying p_w.

through the shared wired bottleneck. The wireless nodes are connected to the ingress router of the wired shared bottleneck via the wireless access links of these heterogeneous wireless architectures. The wired nodes are connected to the same bottleneck via wired links. The link capacities and data packet size are assumed to be 2 Mb/s and 1 KB, respectively. The one-way delay of the wired bottleneck link is assumed to be 10 ms. Each data packet is assumed to be acknowledged by a single ACK, i.e., $b = 1$. The packets are lost in the wireless link with the packet loss probability of p_w. The sliding window duration for the measurement of p_w is set to be $10{\cdot}\mathrm{RTT}$. The investigation of the effects of these parameters on the protocol performance are left for future study. All simulations are performed for a duration of 600 s. For low delay environments such as WLAN, Snoop [3] is proven to be an efficient TCP-aware link-layer solution that significantly increases the TCP performance. WTCP [25] is also known to achieve the highest throughput performance in WWAN environments. To the best of our knowledge, TCP-Westwood [6] is shown to significantly improve the throughput performance in the hybrid wireless/wired networks. Similarly, TCP-Peach+ [1], [2] is shown to be the best reliable transport protocol for satellite networks. Therefore, in order to show that TCP-ATL attains very high performance in all of the heterogeneous wireless environments as shown in Fig. 5, its performance is compared with the performance of Snoop (with TCP-Sack) [3] and TCP-Westwood [6] for WLAN; WTCP [25] and TCP-Westwood [6] for wide-area 3G cellular; TCP-Peach+ [1], [2] and TCP-Westwood [6] for satellite environments.

- **WLAN**: The results of three set of experiments for $d_w = 10, 50$, and 150 ms and varying p_w are plotted in Figs. 6–8, respectively. As shown in Fig. 6, both TCP-ATL and TCP-Westwood outperform Snoop for low packet loss rates, i.e., $p_w < 0.005$ and $d_w = 10$ ms. As p_w increases, $p_w \geq 0.005$, and TCP-Westwood experiences much more significant performance degradation compared with TCP-ATL and Snoop (with TCP-Sack). This is because the adaptive congestion control of TCP-ATL helps to compensate the throughput degradation by dynamically adjusting its AIMD parameters as explained in Section IV-A. Snoop also achieves high performance even in such high packet loss rates because of its local retransmission strategy. Nevertheless, it does not compensate for the effects of the additional wireless access delay on the throughput efficiency. Moreover, in the Snoop implementation, the additional buffering performed at the access point in order to snoop the end-to-end connection brings additional increase in the access delay. This further amplifies the throughput performance difference between the Snoop and TCP-ATL which can handle the adverse effects of the wireless access delay as explained in Section III-B. For a typical WLAN environment with a wide range of p_w, TCP-ATL significantly improves throughput over Snoop and TCP-Westwood. For $p_w = 5.10^{-2}$, TCP-ATL outperforms TCP-Westwood approximately by 234% and Snoop by 47%.

- **3G Cellular**: For this environment, we also compare the performance of TCP-ATL with WTCP in addition to TCP-Westwood. As shown in Fig. 7, TCP-ATL outperforms WTCP for $d_w = 50$ ms and for all p_w values from 10^{-6} to 10^{-1}. However, WTCP achieves higher throughput performance than TCP-Westwood. As the wireless link delay increases, i.e., $d_w = 50$ ms, the throughput improvement by TCP-ATL over TCP-Westwood further increases as shown in Fig. 7. This is mainly because TCP-ATL performance does not degrade with increasing d_w, while TCP-Westwood's performance is significantly affected. This is because the bandwidth estimation based on the ACK reception rate used by the TCP-Westwood starts to lose accuracy because of late arriving feedback [6]. Furthermore, WTCP also outperforms TCP-Westwood in this environment due to the same reason. For $p_w = 10^{-3}$, TCP-Westwood experiences 28% throughput degradation due to the increase in d_w from 10 to 50 ms, while TCP-ATL performance almost remains unaffected. For $p_w = 5.10^{-2}$ and $d_w = 50$ ms, TCP-ATL outperforms WTCP with 91% and TCP-Westwood with approximately 468% throughput improvement.
- **Satellite**: As shown in Fig. 8, TCP-ATL throughput does not significantly degrade in the typical satellite environment, i.e., $d_w = 150$ ms, compared with the previous scenarios. Recall that $d_w = 150$ ms is the one-way wireless link delay. For $p_w = 10^{-3}$ and $d_w = 150$ ms, the TCP-ATL performance decreases by only 6% compared with $d_w = 50$ ms, while TCP-Westwood experiences 48% throughput degradation. This is again because of the adaptive congestion control scheme of the TCP-ATL, which updates its AIMD parameters according to the p_w and d_w. TCP-ATL throughput also shows similar behavior for varying p_w as shown in Fig. 8. However, the throughput degradation due to high p_w is not as severe as it is for TCP-Westwood and TCP-Peach+. TCP-ATL improves the throughput up to 7.6 times over TCP-Westwood for $d_w = 150$ ms. Note also that TCP-ATL outperforms TCP-Peach+ by 125% throughput improvement for $p_w = 10^{-2}$ and $d_w = 150$ ms.

Consequently, instead of using different transport protocols developed for specific architectures, TCP-ATL achieves high-throughput performance in all of the three different wireless architectures by adapting its protocol configuration, as shown in Fig. 5.

2) Roaming Between Heterogeneous Architectures: We also perform simulation experiments to evaluate the throughput achieved by the TCP-ATL, while the mobile terminal roams between different wireless architectures. Here, the mobile terminal performs vertical handoff between three different wireless architectures (WLAN, 3G, satellite), i.e., N_1, N_2, and N_3, at $t = t_1$ and $t = t_2$ during its connection period. The simulations are performed for a duration of 600 s. $p_w = 10^{-5}$, 10^{-4}, 10^{-3}, and $d_w = 10$, 50, and 150 ms for WLAN, 3G cellular, and satellite environments, respectively. The other simulation parameters are the same as the ones used in Section V-A1. The simulations are performed for various (t_1, t_2) and (N_1, N_2, N_3) combinations. Furthermore, we perform the same simulation experiments for two different cases.

1) **Wireless Source**: In this case, the wireless mobile terminal is the source and roaming between heterogeneous wireless architectures of the NGWI, while it has an active connection. For example, in the first experiment, the connection starts in WLAN at $t = 0$, and then at $t_1 = 100$ s the source migrates from WLAN to 3G and at $t_2 = 300$ s from 3G to the satellite environment. As shown in Table II, TCP-ATL maintains its high-throughput performance (around 1.9 Mb/s), regardless of what portion of the connection takes place in which wireless architecture.
2) **Wireless Receiver**: In this scenario, where the wireless mobile terminal is the receiver, the throughput achieved by the TCP-ATL does not vary significantly as shown in Table III. Although there exists very small throughput degradation, this is mainly due to loss of the packets carrying p_w and d_w from the wireless receiver node to the sender. In this case, TCP-ATL source cannot immediately respond to the variation in the p_w and d_w and, hence, throughput performance is affected. Furthermore, another reason for this throughput degradation is slightly delayed congestion control adaptation caused by the additional delay incurred, while the wireless receiver forwardsf the wireless link parameters, i.e., p_w and d_w, to the source as discussed in *case 2* of the detailed protocol operation presented in Section IV-A. Despite such a slight variation, the achieved throughput is still around 1.9 Mb/s for all of the experiment scenarios, as shown in Table III.

TABLE II
THROUGHPUT ACHIEVED WHILE (WIRELESS SOURCE) ROAMING BETWEEN HETEROGENEOUS WIRELESS ARCHITECTURES

Exp. #	$(t_1; t_2)$	$(N_1; N_2; N_3)$	***Tput (Mb/s)***
1	(100,300)	(WLAN,3G,Sat)	1.8870
2	(200,300)	(WLAN,Sat,3G)	1.9149
3	(200,500)	(Sat,WLAN,3G)	1.8919
4	(200,400)	(Sat,3G,WLAN)	1.8979
5	(100,200)	(3G,WLAN,Sat)	1.8701
6	(400,500)	(3G,Sat,WLAN)	1.9210

TABLE III
THROUGHPUT ACHIEVED WHILE (WIRELESS RECEIVER) ROAMING BETWEEN HETEROGENEOUS WIRELESS ARCHITECTURES

Exp. #	$(t_1; t_2)$	$(N_1; N_2; N_3)$	***Tput (Mb/s)***
1	(100,300)	(WLAN,3G,Sat)	1.8521
2	(200,300)	(WLAN,Sat,3G)	1.8943
3	(200,500)	(Sat,WLAN,3G)	1.8755
4	(200,400)	(Sat,3G,WLAN)	1.8771
5	(100,200)	(3G,WLAN,Sat)	1.8561
6	(400,500)	(3G,Sat,WLAN)	1.9093

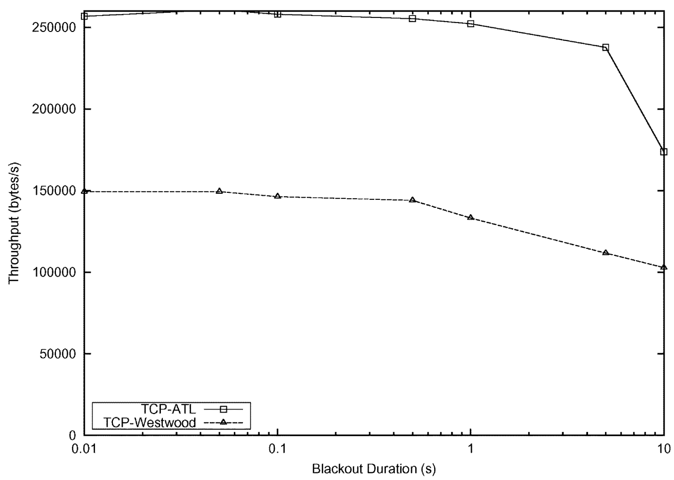

Fig. 9. Throughput comparison of TCP-ATL and TCP-Westwood for varying blackout durations for $d_w = 50$ ms and $p_w = 10^{-3}$.

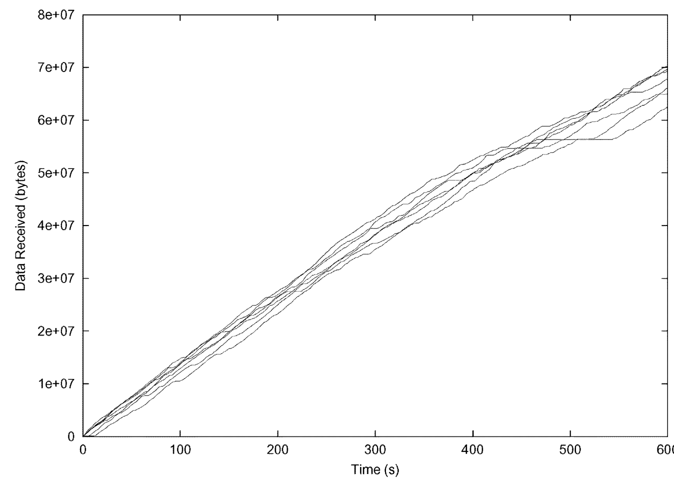

Fig. 10. Intraprotocol fairness: Total bytes received by each of the TCP-ATL sinks with time for $d_w = 10$ ms and $p_w = 10^{-3}$.

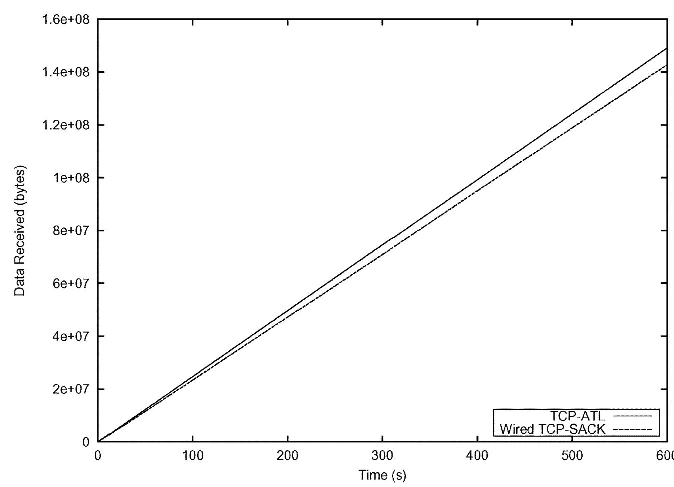

Fig. 11. Fairness to wired TCP flows: Total bytes received by the TCP-ATL and wired TCP sinks with time for $d_w = 10$ ms and $p_w = 10^{-3}$.

Therefore, regardless of whether the wireless mobile terminal is the source or the receiver end of the end-to-end connection, TCP-ATL can provide seamless high network utilization, while mobile users roam between heterogeneous wireless architectures as shown in Fig. 5.

3) Blackout: The throughput achieved by TCP-ATL and TCP-Westwood for varying blackout durations between 10 ms and 10 s are given in Fig. 9. Here, we assume $d_w = 50$ ms and $p_w = 10^{-3}$. The simulations are performed for a duration of 600 s, where the blackout occurs at $t = 300$ s.

As shown in Fig. 9, the throughput decreases with increasing blackout duration, as expected. This decrease is observed in both of the curves representing the TCP-ATL and TCP-Westwood. TCP-ATL outperforms TCP-Westwood for all blackout duration values. This is achieved by the dynamically adaptive congestion control used by TCP-ATL. When a blackout occurs due to mobility or fading by environmental obscurations, this immediately increases p_w. The TCP-ATL source then calculates new α using (9) to adapt the congestion control configuration to the new link conditions. Therefore, TCP-ATL can quickly recover from the throughput degradation due to a blackout as explained in Section IV-A. For the blackout duration of 1 s, TCP-ATL achieves 89% throughput improvement over TCP-Westwood.

4) Fairness: For the fairness performance of TCP-ATL, two different scenarios, i.e., intraprotocol fairness and fairness to wired TCP flows, are considered. For the former case, the fairness among TCP-ATL flows sharing the same bottleneck is investigated. For the latter case, the fairness of TCP-ATL to the wired TCP sources, which are sharing the same bottleneck, is explored.

- **Intraprotocol fairness**: In this case, six TCP-ATL sources share the same bottleneck with a capacity of 6 Mb/s for $d_w = 10$ ms and $p_w = 10^{-3}$. In Fig. 10, the data received by each TCP-ATL sink is shown as a function of time. It is observed from Fig. 10 that all sinks receive almost the same amount of data at any point in time during the connection period. Hence, the TCP-ATL sources are all given a fair share of the network resources.
- **Fairness to wired TCP flows**: In this case, five TCP-ATL and five wired TCP-SACK sources share the same bottleneck with capacity 4 Mb/s. TCP-ATL sources are connected to the bottleneck via a wireless link with $d_w = 10$ ms and $p_w = 10^{-3}$. TCP-SACK sources are connected to the bottleneck via wired links, hence, they are wired TCP sourcs.

The total data received by the TCP-ATL and wired TCP sinks are shown in Fig. 11 as a function of time. Let $\Phi(t)$ be the ratio of the total data received by TCP-ATL and wired TCP sinks at time t. Here, although the TCP-ATL sink receives slightly higher amount of data than the wired TCP source, $\Phi(t) \approx 1$ throughout the connection period. This is mainly because of the inherent fairness of the TCP-ATL adaptive congestion control as explained in Section III-B. Therefore, TCP-ATL significantly improves link utilization efficiency, while preserving the overall fairness to the wired TCP sources sharing the same bottleneck.

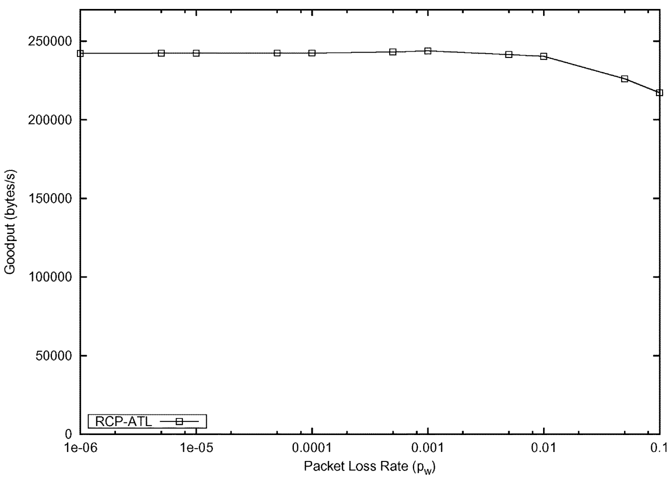

Fig. 12. Goodput achieved by RCP-ATL for $d_w = 10$ ms and varying p_w.

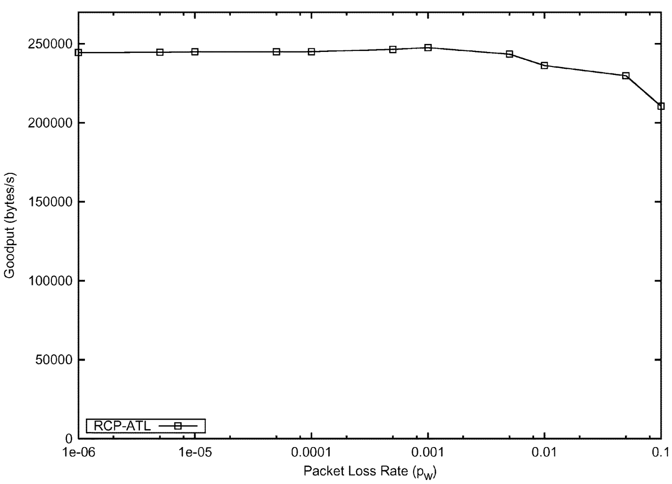

Fig. 13. Goodput achieved by RCP-ATL for $d_w = 50$ ms and varying p_w.

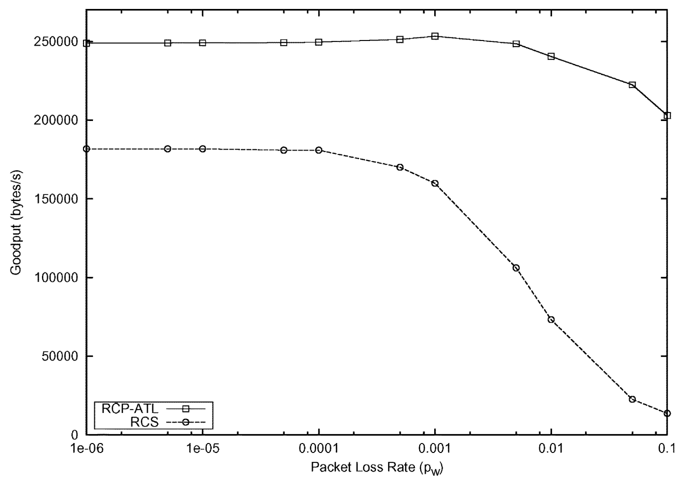

Fig. 14. Goodput comparison of RCP-ATL and RCS for $d_w = 150$ ms and varying p_w.

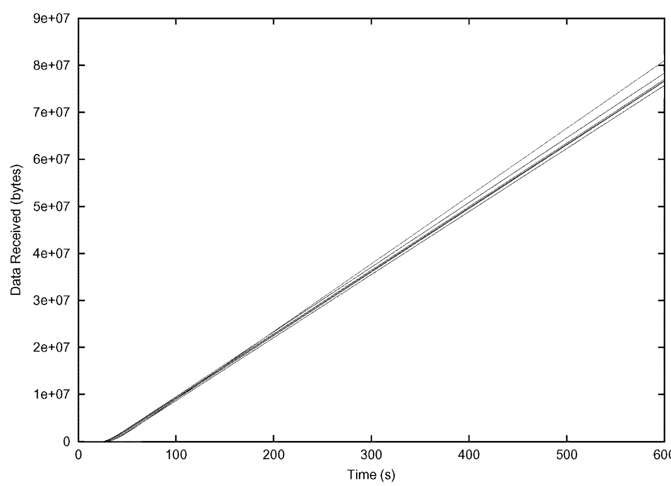

Fig. 15. Intraprotocol fairness: Total bytes received by each of the RCP-ATL sinks with time for $d_w = 10$ ms and $p_w = 10^{-3}$.

B. RCP-ATL Performance

1) Goodput Performance: RCP-ATL goodput performance experiments are also performed for varying p_w between 10^{-6} and 10^{-1} and for $d_w = 10, 50,$ and 150 ms, which represent typical one-way wireless link delay values in WLAN, 3G and satellite networks, respectively. The simulation environment and parameters used are the same with TCP-ATL experiments presented in Section V-A1. To the best of our knowledge, RAP [22] is the best rate control protocol for wired environments. However, it is shown in [27] that RAP has very poor performance in the wireless domain. To our knowledge, however, there is no rate control protocol at the transport layer specifically developed for WLAN and 3G cellular environments. On the other hand, RCS [27] is the rate control scheme that has the highest performance in satellite networks. Therefore, the RCP-ATL goodput performance is compared with only RCS for satellite networks.

For low p_w values, i.e., $p_w < 10^{-2}$ with $d_w = 10$ ms representing WLAN environments, RCP-ATL maintains its goodput performance, as shown in Fig. 12. As p_w increases, RCP-ATL starts to experience a slight performance degradation. As the wireless link delay increases, i.e., $d_w = 50$ ms, the goodput performance of RCP-ATL remains unaffected by the increasing delay, as shown in Fig. 13. This is achieved because of its adaptive rate control explained in Section III-C.

In the typical satellite environment, i.e., $d_w = 150$ ms, we also compare the performance of RCP-ATL with RCS. In this case, RCP-ATL can maintain its high goodput performance over wide range of p_w, while RCS experiences significant degradation for $p_w > 10^{-3}$ as shown in Fig. 14. For $p_w = 10^{-2}$ and $d_w = 150$ ms, RCP-ATL improves goodput over RCS approximately by 228%. Therefore, RCP-ATL can achieve very high goodput performance in all different wireless environments as its TCP counterpart as shown in Fig. 5, where α is now calculated by (13).

2) Fairness: Similar to the Section V-A4, two different scenarios, i.e., intraprotocol fairness and fairness to wired TCP flows, are considered for the fairness performance of RCP-ATL. In the intraprotocol fairness scenario, six RCP-ATL sources share the same bottleneck with capacity of 6 Mb/s for $d_w = 10$ ms and $p_w = 10^{-3}$. As shown in Fig. 15, the data received by each RCP-ATL sink is almost the same at any point in time during the connection period. Hence, the RCP-ATL sources are all given a fair share of the network resources.

To explore the RCP-ATL fairness to the wired TCP flows, the same scenario as in Section V-A4 is used. RCP-ATL and wired TCP-SACK sources share the same bottleneck with a capacity of 4 Mb/s. The data received by the RCP-ATL and the wired

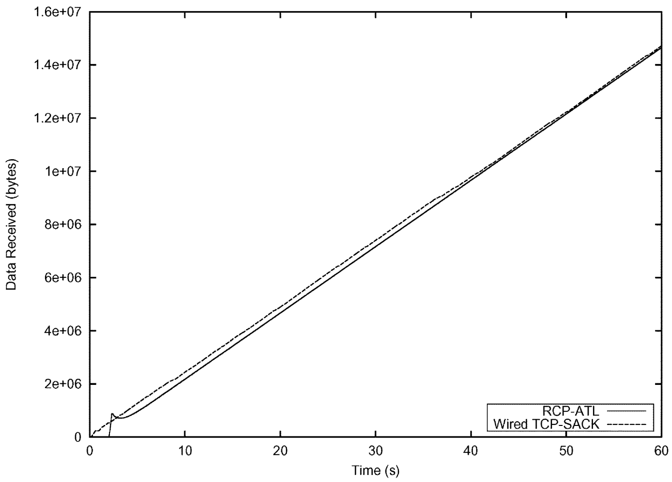

Fig. 16. Fairness to wired TCP flows: Total bytes received by the RCP-ATL and wired TCP sinks with time for $d_w = 10$ ms and $p_w = 10^{-3}$.

TCP sinks are shown in Fig. 16 as a function of time. Let $\Phi(t)$ be the ratio of the data received by RCP-ATL and wired TCP sinks at time t. As shown in Fig. 16, $\Phi(t) \approx 1$ throughout the connection period. Therefore, RCP-ATL maintains fairness to the wired TCP sources sharing the same bottleneck. Note that this fairness is achieved because of the adaptive rate control designed by taking fairness into consideration, as explained in Section III-C.

3) Jitter Issues: The jitter experienced by RCP-ATL and RAP flows are also investigated in terms of the multimedia support performance of the RCP-ATL. Although RAP [22] is mainly developed for wired networks, we compared RCP-ATL with RAP in order to assess the improvement achieved by RCP-ATL due to its adaptivity to varying link conditions. The jitter experienced by each packet transmitted by the RCP-ATL and RAP sources are shown in Fig. 17. In the bottom plot of Fig. 17, it is seen that the packets sent by the RAP source experience much higher jitter than the packets sent by RCP-ATL. This is because RAP halves its data rate at packet losses due to link errors, hence resulting in high delay variation. In order to comply with the instantaneous TCP-friendliness, RCP-ATL also performs rate decrease at each packet loss. However, since it uses higher multiplicative decrease parameter β, RCP-ATL experiences less abrupt rate variation in case of a packet loss than RAP. Therefore, RAP source experiences higher jitter than RAP source. Such improvement in jitter achieved by RCP-ATL is of significant importance for the delay sensitive real-time traffic, where the delay variation directly affects the QoS performance.

VI. Conclusion

The transport layer protocols developed for different wireless architectures do not provide a single solution to address the heterogeneities posed by the NGWI. The inadequacies of the current transport layer solutions necessitate a new transport layer that can address these heterogeneities. To the best of our knowledge, there has been no single proposed transport layer solution to serve for the NGWI objective.

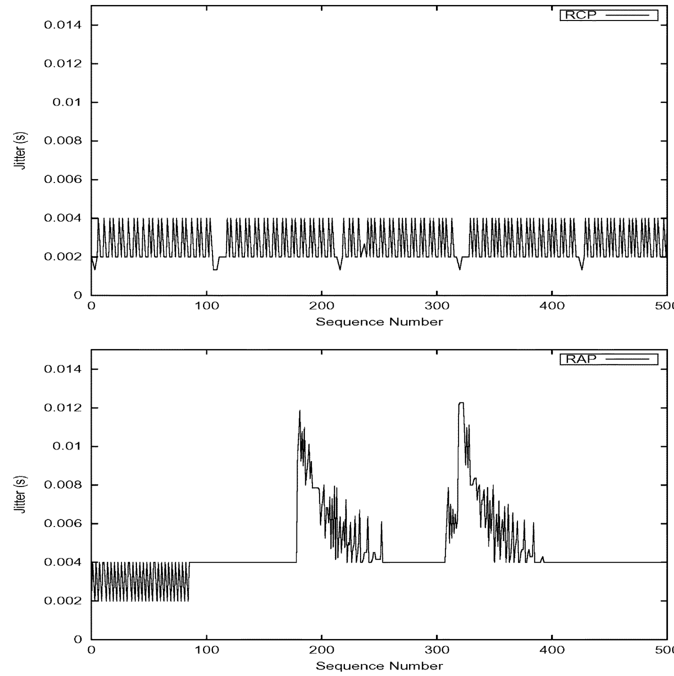

Fig. 17. Jitter experienced by the RCP-ATL and RAP packets for $d_w = 10$ ms and $p_w = 10^{-4}$.

In this paper, a new unified ATL is presented to realize the NGWI objective. ATL incorporates a new adaptive transport protocol (TCP-ATL) for reliable data transport; and a new adaptive AIMD rate control protocol (RCP-ATL) for multimedia delivery in the NGWI. Both TCP-ATL and RCP-ATL deploy new adaptive congestion/rate control methods that dynamically adapt their AIMD control parameters to the current wireless link conditions to maintain high-throughput performance and to provide multimedia traffic support in NGWI. The experiment results showed that ATL protocols can maintain high-performance throughout the different wireless architectures. Their adaptive AIMD configuration methods also help avoid performance degradation for high packet loss rates. Furthermore, since the developed AIMD parameter adjustment strategy takes fairness into consideration, ATL protocols also maintain fairness to the wired TCP sources sharing the same bottleneck. Moreover, ATL protocols do not bring any additional overhead and can be developed on top of any of the existing TCP or AIMD rate control protocols.

Note also that the unified ATL solution avoids implementing different protocol stacks both at the wired and wireless terminals for different wireless architectures, i.e., WLAN, 3G cellular, and satellite. In fact, since the memory efficiency is an important concern for small and resource-constrained wireless devices and it is foreseen that memory will be mostly consumed by the applications, a set of software tools is being developed specifically for protocol and algorithm development with minimum memory footprint for wireless terminals [30]. Hence, instead of separate protocols for different architectures, the single unified ATL implementation utilizes up to three times lower memory footprint size, which is a significant memory saving in terms of software design for memory constrained wireless terminals [30], [31].

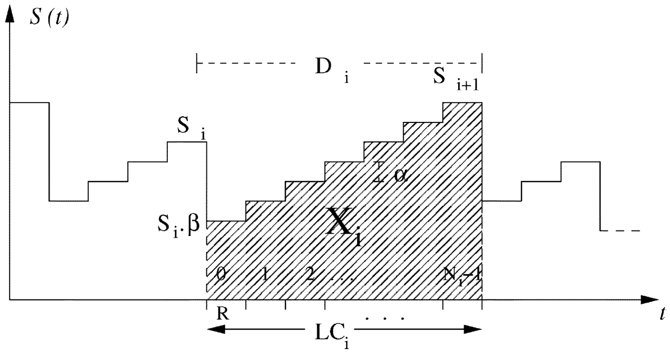

Fig. 18. Data rate change of rate-base general AIMD congestion control.

Furthermore, implementing three different transport protocols would also require a separate mechanism to efficiently hand over the active transport connections between different protocols as the mobile terminal roams between wireless architectures. This involves in per-connection state maintenance, buffering, switching, and synchronization, which lead to significant memory, processing, and software overhead for resource-constrained wireless terminals [32]. Hence, ATL also avoids the need for such a mechanism by providing a single unified transport layer solution which achieves overhead-free and seamless transport connection as the mobile terminal roams across the NGWI.

The experiment results from the experiments demonstrated that both ATL protocols achieve very high performance in the heterogeneous wireless architectures for wide ranges of packet loss probabilities and link delays. TCP-ATL is shown to significantly improve the throughput performance over Snoop (with TCP-Sack) [3] and TCP-Westwood [6] for WLAN environments, WTCP [25] and TCP-Westwood [6] for wide-area 3G cellular, and TCP-Peach+ [1], [2] and TCP-Westwood [6] for satellite environments. Hence, instead of using different transport protocols developed for specific architectures, TCP-ATL achieves high-throughput performance in all of the three different wireless architectures by adapting its protocol configuration. TCP-ATL are also shown to address the blackout situations and maintain its throughput performance. Both TCP-ATL and RCP-ATL is shown to preserve fairness to the wired TCP sources sharing the same bottleneck. RCP-ATL is also shown to improve jitter performance in wireless environments. As a result, the ATL protocol suite addresses the challenges posed by the NGWI and significantly improves the performance for reliable data and multimedia transport in NGWI.

APPENDIX

The time-dependency of the transmission rate $S(t)$ is shown in Fig. 18. The rate-based generic AIMD scheme is assumed to increase the rate S additively with α at each RTT, i.e., $S = S + \alpha$. It throttles the transmission rate, S, multiplicatively by β if a packet loss is detected, i.e., $S = \beta \cdot S$. The steady-state throughput of such *rate-based* AIMD scheme is derived here.

Let $X(t)$ be the total number of packets transmitted in $[0, t)$, which can be calculated by

$$X(t) = \int_0^t S(\tau) \mathrm{d}\tau. \tag{18}$$

If $T(t) = X(t)/t$ is the throughput achieved in $[0, t)$, the steady-state throughput achieved by the rate-based generic AIMD scheme is given by

$$T = \lim_{t \to \infty} \frac{1}{t} \int_0^t S(\tau) \mathrm{d}\tau. \tag{19}$$

Let X_i be the number of packets transmitted in the ith loss cycle LC_i, which starts and ends with rate halving due to congestion decision. If the duration of LC_i is D_i, then the throughput achieved in LC_i is given by $T_i = X_i/D_i$. Assuming the evolution of the transmission rate S_i to be Markov regenerative process with rewards X_i [16], then the steady-state throughput can be calculated by

$$T = \frac{E[X]}{E[D]} \tag{20}$$

where $E[X]$ and $E[D]$ are the means of X_i and D_i, respectively. The transmission rate change during LC_i, i.e., the value of the transmission rate at kth RTT of the LC_i, is given by

$$S_{i,k} = S_i \cdot \beta + k \cdot \alpha. \tag{21}$$

Given that N_i is the total number of RTTs in LC_i, i.e., the rate is throttled in $(N_i - 1)$th RTT, then the transmission rate at the end of the LC_i is expressed by $S_{i+1} = S_i \cdot \beta + N_i \cdot \alpha$. Hence, the expectation of the i.i.d random variable S denoting the transmission rate, i.e., $E[S]$, can be calculated as

$$E[S] = \frac{\alpha}{1 - \beta} \cdot E[N]. \tag{22}$$

If LC_i lasts for $D_i = N_i \cdot R$, where R is the RTT, then the total number of packets transmitted during LC_i can be calculated by

$$X_i = \int_0^{D_i} S_i(t) \cdot \mathrm{d}t. \tag{23}$$

For rate-based AIMD scheme, whose rate change is performed with RTT granularity, this can be calculated by replacing the integration in (23) with the discrete summation and substituting (21) as follows:

$$\begin{aligned} X_i &= \sum_{k=0}^{N_i - 1} S_{i,k} \cdot R \\ &= \sum_{k=0}^{N_i - 1} (S_i \cdot \beta + k \cdot \alpha) \cdot R \\ &= \frac{N_i}{2} [2 \cdot \beta \cdot S_i + \alpha \cdot (N_i - 1)] \cdot R. \end{aligned} \tag{24}$$

Thus, for mutually independent random variables of N and S, the mean of X can be calculated by taking expectation of both sides of (24) and substituting (22) as follows:

$$E[X] = \frac{E[N]}{2} \cdot \left[\left(\frac{1 + \beta}{1 - \beta} \right) \cdot E[N] - 1 \right] \cdot \alpha \cdot R. \tag{25}$$

On the other hand, the total number of packets transmitted in loss cycle i can also be calculated by $n_i + S_{i,N_i - 1} \cdot R$, where n_i and $S_{i,N_i - 1} \cdot R$ are the number of packets transmitted until the dropped packet and in the last RTT, respectively. Therefore, $E[X]$ can also be calculated by

$$E[X] = E[n] + E[S] \cdot R \tag{26}$$

where the expectation of the random variable n is given by

$$\begin{aligned} E[n] &= \sum_k kP[n_i = k] \\ &= \sum_{k=0}^{\infty} k(1-p)^{k-1}p = \frac{1}{p} \end{aligned} \tag{27}$$

where p is the packet loss probability, if loss-based congestion detection is used by the generic AIMD rate-based congestion control algorithm. By substituting (22) and (27) into (26), and equating it to (25), we solve for $E[N]$ and obtain it as follows

$$E[N] = \frac{3-\beta}{2(1+\beta)} \left[1 + \sqrt{1 + \frac{8(1-\beta^2)}{\alpha R p (3-\beta)^2}} \right]. \tag{28}$$

Thus, it follows from (20), (25) and (28) that the steady-state throughput of the rate-based general AIMD rate control scheme as a function of rate-increase and decrease parameters, i.e., α and β, and RTT R and the packet loss probability p, can be expressed as follows:

$$\mathcal{T}^r_{\alpha,\beta}(p,R) = \frac{\alpha}{4(1-\beta)} \left[1 + \beta + \sqrt{(3-\beta)^2 + \frac{8(1-\beta^2)}{\alpha R p}} \right]. \tag{29}$$

REFERENCES

[1] I. F. Akyildiz, G. Morabito, and S. Palazzo, "TCP-Peach: A new congestion control scheme for satellite IP networks," *IEEE/ACM Trans. Networking*, vol. 9, pp. 307–321, June 2001.

[2] I. F. Akyildiz, X. Zhang, and J. Fang, "TCP-Peach+: Enhancement of TCP-Peach for satellite IP networks," *IEEE Commun. Lett.*, vol. 6, pp. 303–305, July 2002.

[3] H. Balakrishnan, S. Seshan, E. Amir, and R. H. Katz, "Improving TCP/IP performance over wireless networks," in *Proc. ACM MOBICOM*, Nov. 1995, pp. 2–11.

[4] H. Balakrishnan, V. N. Padmanabhan, S. Seshan, and R. H. Katz, "A comparison of mechanisms for improving TCP performance over wireless links," *IEEE/ACM Trans. Networking*, vol. 5, pp. 756–769, Dec. 1997.

[5] H. Balakrishnan, H. S. Rahul, and S. Seshan, "An integrated congestion management architecture for Internet hosts," in *Proc. ACM SIGCOMM*, Sept. 1999, pp. 175–187.

[6] C. Casetti, M. Gerla, S. Mascolo, M. Y. Sanadidi, and R. Wang, "TCP-Westwood: Bandwidth estimation for enhanced transport over wireless links," in *Proc. ACM MOBICOM*, Rome, July 2001, pp. 287–297.

[7] S. Cen, C. Pu, and J. Walpole, "Flow and congestion control for Internet media streaming applications," in *Proc. SPIE Multimedia Computing Networking*, Jan. 1998, pp. 250–264.

[8] S. Floyd, M. Handley, J. Padhye, and J. Widmer, "Equation-based congestion control for unicast applications," in *Proc. ACM SIGCOMM*, Aug. 2000, pp. 45–58.

[9] S. Floyd, M. Handley, and J. Padhye. (2000, May) A comparison of equation-based and aimd congestion control. ACIRI, Tech. Rep. [Online]. Available: http://www.aciri.org/tfrc/aimd.pdf

[10] T. R. Henderson and R. H. Katz, "Transport protocols for Internet-compatible satellite networks," *IEEE J. Select. Areas Commun.*, vol. 17, pp. 326–344, Feb. 1999.

[11] H.-Y. Hsieh and R. Sivakumar, "A transport layer approach for achieving aggregate bandwidths on multi-homed mobile hosts," in *Proc. ACM MOBICOM*, Atlanta, GA, Sept. 2002, pp. 83–94.

[12] IEEE 802.11, Wireless LAN medium access control (MAC) and physical layer (Phy) specification, 1999.

[13] H. Inamura, G. Montenegro, R. Ludwig, A. Gurtov, and F. Khafizov, "TCP over second (2.5G) and third (3G) generation wireless networks," IETF, Internet-Draft draft-ietf-pilc-2.5g3g-12, Dec. 2002.

[14] T. V. Lakshman and U. Madhow, "The performance of TCP/IP for networks with high bandwidth-delay products and random loss," *IEEE/ACM Trans. Networking*, vol. 5, pp. 336–350, June 1997.

[15] J. R. Li, S. Ha, and V. Bharghavan, "HPF: A transport protocol for supporting heterogeneous packet flows in the Internet," in *Proc. IEEE INFOCOM*, vol. 3, Mar. 1999, pp. 543–550.

[16] D. Logothetis, K. S. Trivedi, and A. Puliafito, "Markov regenerative models," in *Proc. Int. Computer Performance Dependability Symp.*, Erlangen, Germany, 1995, pp. 134–143.

[17] R. Ludwig and R. H. Katz, "The Eifel algorithm: Making TCP robust against spurious retransmissions," *ACM Comput. Commun. Rev.*, vol. 30, no. 1, pp. 30–36, Jan. 2000.

[18] M. Mathis, J. Mahdavi, S. Floyd, and A. Romanov, "TCP selective acknowledgment options," IETF, RFC 2018, 1996.

[19] The Network Simulator, ns-2 [Online]. Available: http://www.isi.edu/nsnam/ns/index.html

[20] J. Padhye, V. Firoio, D. Towsley, and J. Kurose, "Modeling TCP-Reno performance: A simple model and its empirical validation," *IEEE/ACM Trans. Networking*, vol. 8, pp. 133–145, Apr. 2000.

[21] R. Puri, K.-W. Lee, K. Ramchandran, and V. Bharghavan, "An integrated source transcoding and congestion control paradigm for video streaming in the Internet," *IEEE Trans. Multimedia*, vol. 3, pp. 18–32, Mar. 2001.

[22] R. Rejaie, M. Handley, and D. Estrin, "RAP: An end-to-end rate-based congestion control mechanism for realtime streams in the Internet," in *Proc. IEEE INFOCOM*, vol. 3, Mar. 1999, pp. 1337–1345.

[23] O. Sallent, J. P. Romero, R. Agusti, and F. Casadevall, "Provisioning multimedia wireless networking for better QoS: RRM strategies for 3G W-CDMA," *IEEE Commun. Mag.*, vol. 41, pp. 100–106, Feb. 2003.

[24] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson, "RTP: A transport for real-time applications," IETF, RFC 1889, Jan. 1996.

[25] P. Sinha, N. Venkitaraman, R. Sivakumar, and V. Bharghavan, "WTCP: A reliable transport protocol for wireless wide-area networks," in *Proc. ACM MOBICOM*, Seattle, WA, Aug. 1999, pp. 231–241.

[26] W. Stevens, "TCP Slow start, congestion avoidance, fast retransmit, and fast recovery algorithms," IETF, RFC 2001.

[27] J. Tang, G. Morabito, I. F. Akyildiz, and M. Johnson, "RCS: A rate control scheme for real-time traffic in networks with high bandwidth-delay products and high bit error rates," in *Proc. IEEE INFOCOM*, vol. 1, Apr. 2001, pp. 114–122.

[28] Y. R. Yang and S. S. Lam, "General AIMD congestion control," in *Proc. IEEE ICNP*, Osaka, Japan, Nov. 2000, pp. 187–198.

[29] F. Yang, Q. Zhang, W. Zhu, and Y.-Q. Zhang, "An end-to-end TCP-friendly streaming protocol for multimedia over wireless internet," in *Proc. IEEE ICME 2003*, vol. 2, July, 6–9 2003, pp. 429–432.

[30] Sun Microsystems. The connected limited device configuration (CLDC) hotspot implementation virtual machine (JAVA 2 Platform, Micro Edition). Sun Microsystems White Paper. [Online]. Available: http://java.sun.com/products/cldc/

[31] K. Laraqui, "Bridge-on-chip: The enabler of next-generation mobile internet and voice-over-IP," *IEEE Circuits Devices Mag.*, vol. 18, pp. 7–12, Mar. 2002.

[32] F. Sultan, K. Srinivasan, D. Iyer, and L. Iftode, "Migratory TCP: Connection migration for service continuity in the Internet," in *Proc. IEEE ICDCS*, Vienna, Austria, July 2002.

Özgür B. Akan (S'02) received the B.Sc. and M.Sc. degrees in electrical and electronics engineering from Bilkent University and Middle East Technical University, Ankara, Turkey, in 1999 and 2001, respectively. He is currently working toward the Ph.D. degree at the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta.

He is a Research Assistant in the Broadband and Wireless Networking Laboratory, School of Electrical and Computer Engineering, Georgia Institute of Technology. His current research interests include adaptive transport protocols for heterogeneous wireless architectures, next-generation wireless networks, and deep-space communication networks.

Ian F. Akyildiz (M'86–SM'89–F'96) received the B.S., M.S., and Ph.D. degrees in computer engineering from the University of Erlangen, Nuernberg, Germany, in 1978, 1981, and 1984, respectively.

Currently, he is the Ken Byers Distinguished Chair Professor with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, and Director of Broadband and Wireless Networking Laboratory. He is the Editor-in-Chief of *Computer Networks* (Elsevier) and the founding Editor-in-Chief of *Ad Hoc Networks Journal* (Elsevier). His current research interests are in sensor networks, interplanetary Internet, wireless networks, satellite networks and the next-generation Internet.

Dr. Akyildiz is an Association for Computing Machinery (ACM) Fellow (1996). He received the "Don Federico Santa Maria Medal" for his services to the Universidad of Federico Santa Maria, Valparaíso, Chile, in 1986. He served as a National Lecturer for ACM from 1989 until 1998 and received the ACM Outstanding Distinguished Lecturer Award for 1994. He received the 1997 IEEE Leonard G. Abraham Prize Award (IEEE Communications Society) for his paper entitled "Multimedia Group Synchronization Protocols for Integrated Services Architectures" published in the IEEE JOURNAL OF SELECTED AREAS IN COMMUNICATIONS (JSAC) in January 1996, the 2002 IEEE Harry M. Goode Memorial Award (IEEE Computer Society) with the citation "for significant and pioneering contributions to advanced architectures and protocols for wireless and satellite networking," the 2003 IEEE Best Tutorial Award (IEEE Communications Society) for his paper entitled "A Survey on Sensor Networks," published in the *IEEE Communication Magazine* in August 2002, and the 2003 ACM Sigmobile Outstanding Contribution Award with the citation: "for pioneering contributions in the area of mobility and resource management for wireless communication networks." He was the Technical Program Chair and General Chair for several IEEE and ACM conferences including IEEE INFOCOM, ACM MOBICOM, and IEEE ICC. He is a Past Editor for the IEEE/ACM TRANSACTIONS ON NETWORKING (1996–2001), as well as for the IEEE TRANSACTIONS ON COMPUTERS (1992–1996).