

Empirical Evaluation of TCP Variants and an Amendment to Veno on a Controlled AWGN-Impaired WiFi Link

^{1st} Luke Friggieri Cordina

Dept. of Communications and Computer Engineering
University of Malta
 Msida, Malta
 luke.friggieri.23@um.edu.mt

^{2nd} Etienne-Victor Depasquale

Dept. of Communications and Computer Engineering
University of Malta
 Msida, Malta
 etienne.depasquale@um.edu.mt

^{3rd} Owen Casha

Dept. of Microelectronics and Nanoelectronics
University of Malta
 Msida, Malta
 owen.casha@um.edu.mt

Abstract—Most comparative evaluations of TCP congestion control over Wi-Fi rely on simulation or emulation and treat the wireless channel as an idealised loss process. This paper reports a fully empirical, hardware-calibrated study of four TCP variants (CUBIC, Westwood, Veno, and Vegas) over a physical 802.11g ad-hoc link inside an anechoic chamber, where receiver SNR was swept from 35 dB to 5 dB using a Vector Network Analyser-calibrated AWGN injection path. Goodput, RTT, and retransmissions were measured across $n = 5$ repeated 55-second single-stream downstream trials per condition and evaluated with one-way ANOVA, Tukey HSD, and Cohen’s d . Results show Vegas achieves markedly lower latency at a small goodput cost, Westwood best sustains goodput at 11 dB via bandwidth estimation, and Veno exhibits a goodput drop at this point consistent with a static threshold that does not scale with modern congestion window sizes. An adaptive β threshold and an RTT Growth Guard, grounded in 802.11g MAC retry timing, are derived and implemented as a loadable kernel module amending TCP Veno, then re-measured on the same calibrated testbed. Results are mixed: Adaptive Beta yields a significant RTT reduction at the clean-channel ceiling and uniquely sustains goodput at the extreme low-SNR floor, but neither amendment improves goodput at 11 dB, and both show significantly higher retransmissions than legacy Veno at 5 dB. These are reported as a partial, honest validation, not a wholesale improvement.

Index Terms—TCP congestion control, Wi-Fi, empirical testbed, AWGN, signal-to-noise ratio, TCP Veno, bufferbloat, anechoic chamber, loadable kernel module

I. INTRODUCTION

TCP remains the dominant transport for reliable data delivery, but many congestion-control algorithms assume packet loss is a reliable congestion signal [1], an assumption that breaks down over Wi-Fi, where interference, fading, attenuation, and MAC-layer retries distort the signals TCP relies on. Loss-based schemes such as CUBIC may interpret wireless-induced loss as congestion, causing unnecessary window reduction and reduced goodput [2]–[4].

End-to-end TCP observes only aggregate symptoms of the path, without directly observing whether they stem from RF impairment, queue build-up, or contention. This motivates isolating the single-sender, SNR-driven component before reintroducing multi-flow effects such as the “TCP Victim” phenomenon under 802.11ax [7]. AWGN in an anechoic chamber is used here as a controlled means of inducing repeatable, channel-originated impairment on real hardware, not as a complete model of indoor Wi-Fi propagation.

Much of the literature is purely simulation-based [5] or relies on emulators over idealised channel models [6], which cannot reproduce the coupling between physical-layer retry timing, buffer occupancy, and RTT inference that governs real TCP behaviour. Empirical work that does exist typically benchmarks a narrow subset of variants or evaluates only throughput while neglecting latency and retransmission dynamics jointly [5], [6], [8].

This paper makes two contributions. First, a controlled empirical baseline using real 802.11g hardware in an RF-isolated anechoic chamber with a VNA-calibrated AWGN path, sweeping true receiver SNR from 35 dB to 5 dB for CUBIC, Westwood, Veno, and Vegas. Second, having diagnosed a scaling limitation in TCP Veno from this baseline, an adaptive threshold and RTT Growth Guard are derived and packaged as a loadable kernel module (LKM) amending only Veno’s classifier, then empirically re-measured on the same calibrated testbed. This distinguishes the present work from the underlying dissertation, where kernel-level validation was future work. We report the follow-up results in full, including SNR regions where the amendment did not improve outcomes.

Section II summarises related work and the four variants under test. Section III details the RF testbed, calibration procedure, and statistical pipeline. Section IV presents baseline results. Section V derives the amendment and reports the LKM validation session. Section VI concludes with limitations and future work.

II. BACKGROUND AND RELATED WORK

Aoyagi et al. [6] used the Mininet-WiFi emulator alongside limited real-testbed measurements to compare CUBIC and BBR under bufferbloat, but did not physically calibrate SNR. Taha et al. [5] compared Reno, Vegas, and Westwood over Wi-Fi and WiMAX purely through simulation. Pang et al. [8] benchmarked Reno, Westwood, and Veno using WLAN emulation plus real hardware, but did not sweep a calibrated physical channel variable. Du et al. [7] conducted genuinely empirical measurements on 802.11ax hardware and identified the “TCP victim” unfairness phenomenon, but this required multi-client contention as a precondition and did not isolate single-flow SNR causality.

No prior study combines a fully empirical, VNA-calibrated physical channel, simultaneous multi-metric capture, formal inferential statistics across a swept SNR range for more than two variants, and an empirically re-measured LKM amendment derived from that same diagnosis. This is the gap addressed here.

A. Algorithmic Mechanisms

CUBIC [2] grows the congestion window as a cubic function of time since the last loss event, using the prior window maximum W_{max} as an inflection point, making growth independent of RTT but treating every loss as a congestion signal.

TCP Westwood [9] estimates available bandwidth (BWE) from ACK arrival rate and sets $ssthresh = (BWE \times RTT_{min})/MSS$ upon loss, distinguishing wireless loss from true congestion.

TCP Veno [10] is a hybrid scheme that computes a backlog estimate $D = cwnd \times (rtt - rtt_{base})/rtt$ and compares it against a static threshold β to decide whether a loss event is congestive or wireless-induced.

TCP Vegas [11] proactively compares $Diff = Expected - Actual$ throughput against thresholds α, β to bound queue occupancy before loss occurs, trading peak goodput for latency control.

III. METHODOLOGY

An anechoic chamber eliminates ambient RF interference. An ad-hoc (IBSS) topology removes access-point scheduling variability, isolating TCP/SNR causality; this single-flow topology means no random-access contention was present during measurement, a scope boundary discussed in Section VI. *iperf3* captures goodput, RTT, and retransmissions simultaneously for a single downstream TCP stream per trial. The identical testbed, calibration, and statistical pipeline were reused without modification for the LKM validation session in Section V.

A. RF Signal Path and Calibration

A 180° hybrid ring coupler was inserted into the receive path (Fig. 1), with an AWGN source at its isolated noise port ($S_{31} = 36.6$ dB isolation, preventing transmitter desensitisation). A variable attenuator enabled repeatable SNR steps from 35 dB to 5 dB. The true incident signal power P_1 was recovered from the coupler’s scattering parameters:

$$P_1 = \frac{P_{\Delta} - |S_{43}|^2 P_3}{|S_{41}|^2} \quad (1)$$

where P_{Δ} is power at the difference port, P_3 is AWGN source power, and S_{41}, S_{43} are measured transfer coefficients. All RF components, including cable insertion loss (TX: 2.58 dB, RX: 2.84 dB at 2.412 GHz), were characterised on a VNA beforehand.

B. Automated Testing Framework and Statistics

A custom framework switched TCP variants via kernel *sysctl*, executing five repeated 55-second *iperf3* single-stream downstream runs per variant per SNR level ($n = 5$ per cell; $df_{between} = 3$, $df_{within} = 16$ for the four-variant baseline; $df_{between} = 2$, $df_{within} = 12$ for the three-variant LKM session). One-way ANOVA tested inter-variant differences at each bin; significant results




Fig. 1. Hardware setup used for the controlled AWGN-impaired Wi-Fi measurements, showing the anechoic-chamber placement of the two endpoints and the external RF chain used to inject receiver-side noise through the hybrid ring coupler.

($p < 0.05$) were followed by Tukey HSD and Cohen's d . Shapiro-Wilk and Levene's tests indicated the non-normality typical of right-skewed, zero-bounded network metrics. ANOVA was retained over non-parametric alternatives because groups were exactly balanced ($n = 5$ throughout), a condition under which the F-test is robust to moderate non-normality, and because significant effect sizes were consistently large ($d > 1.0$), substantially exceeding any skew-induced bias. The full dataset, analysis scripts, and kernel module source code are publicly available [13].

IV. BASELINE RESULTS

The IBSS/802.11g topology enforced a hard goodput ceiling of ≈ 21 Mbps, consistent with established 802.11g MAC bounds [12].

At 35 dB, loss-based variants (CUBIC, Westwood, VenO) converged near 20.8 Mbps, while Vegas showed a significant but small goodput deficit of ≈ 0.8 Mbps ($F = 42.35$, $p < 0.001$; $d > 3.2$). At 11 dB, where 802.11 MAC retries begin propagating to the transport layer, Westwood retained the highest mean




Fig. 2. Mean goodput of CUBIC, Westwood, VenO, and Vegas across the measured SNR sweep.

goodput (13.21 Mbps), CUBIC dropped to 9.20 Mbps ($F = 7.43$, $p < 0.001$), and VenO showed a comparable drop (Fig. 2).

Latency showed the inverse pattern (Fig. 3): at 35 dB, Vegas' RTT (2.99 ms) was significantly lower than the 29–38 ms range for loss-based variants ($d > 3.0$, $p < 0.0001$ at 20 dB, Tukey HSD).

Retransmissions (Fig. 4) remained low until the 5 dB




Fig. 3. Average RTT of the four baseline variants across the measured SNR sweep (log scale).




Fig. 4. Mean TCP-level retransmission counts across the measured SNR sweep.

floor, where CUBIC, Westwood, and Vegas rose above 25 events; Veno remained lower at 5.8 ($F = 4.32$, $p = 0.0106$; $d > 1.4$), suggesting its conservative window growth limited retry-driven loss.

V. VENO AMENDMENT: DERIVATION AND VALIDATION

Veno forgives a loss as non-congestive only if

$$\frac{r_{tt} - r_{tt_{base}}}{r_{tt}} < \frac{\beta}{2 \cdot cwnd} \quad (2)$$

with static kernel constant $\beta = 6$. At 11 dB, $cwnd \approx 67$ packets was estimated from the bandwidth-delay product using measured goodput 11.12 Mbps, RTT 69.62 ms, and $MSS = 1448$ bytes—a steady-state approximation, not a direct kernel-instrumented reading, which we note as a limitation on the precision, though not the direction, of the diagnosis below. This bounds tolerated RTT inflation to $\approx 4.5\%$, below the 10–15% typical of an 802.11g retry cycle [12]. In Veno’s 2003 design context ($cwnd \approx 18$), the same formula tolerated $\approx 16.7\%$, motivating the corrections below.

A. Adaptive Beta Threshold

$$\beta_{a,c} = \max\left(3, \left\lfloor \frac{cwnd}{6} \right\rfloor\right) \quad (3)$$

converted to fixed-point via $\beta_{a,s} = \beta_{a,c} \ll s_V$, where s_V is the Linux macro `V_PARAM_SHIFT` in `net/ipv4/tcp_veno.c`. Divisor 6 is an analytic-continuity choice, not an independently tuned constant: it numerically coincides with legacy β , so $\beta_{a,c} = 3 = \beta/2$ reproduces legacy behaviour exactly at $cwnd \approx 18$, while $\max(3, \cdot)$ prevents regression for smaller windows. This yields an asymptotic tolerance of $\approx 16.7\%$ independent of $cwnd > 18$; the floor, not the divisor, dominates at legacy window sizes. This amendment targets 11 dB specifically.

B. RTT Growth Guard

A second, independent mechanism targets the 35 dB ceiling, where Veno has no loss-driven trigger to halt growth:

$$RTT_{current} > RTT_{min} \times k, \quad k = 2.0 \quad (4)$$

k is grounded in 802.11g MAC retry timing [12]: overhead alone (DIFS = $28 \mu s$, mean CW_{min} backoff $\approx 67.5 \mu s$, ACK $\approx 39 \mu s$) totals only $\approx 135 \mu s$, but automatic rate fallback under the retry (e.g. 54 to 6 Mbps) adds ≈ 1.8 ms of frame airtime for a full 1500-byte frame, together reproducing the observed 2–3 ms single-retry inflation. Against the empirical $RTT_{min} = 2.99$ ms (the Vegas proxy for a near-empty queue), this gives an inflation ratio of ≈ 1.67 –2.0, motivating $k = 2.0$. A sensitivity check at $k = 3.0$ still bounds RTT to 8.97 ms versus the 37.84 ms legacy baseline (76.3% reduction), so the guard’s benefit is not contingent on precise calibration of k ; its observed 35 dB failure (Section V-C) is a separate implementation issue. We note that RTT_{min} was taken from the Vegas baseline and was not independently verified for stability between slow-start and steady state; upward drift during slow-start would itself inflate the guard’s freeze threshold, a plausible contributing factor to the anomaly reported below and an open question requiring instrumentation, not a confirmed cause.

C. Kernel Module Validation Session

Both amendments were implemented as a single LKM registering a modified identifier alongside unmodified TCP Veno, selectable via `sysctl`. Legacy Veno and both amended variants were re-measured on the identical




Fig. 5. Measured goodput for Legacy VenO, Adaptive Beta, and the RTT Growth Guard.




Fig. 6. Measured average RTT for Legacy VenO, Adaptive Beta, and the RTT Growth Guard.

calibrated testbed ($n = 5$ per condition, single down-stream stream per trial).

At 35 dB (Fig. 6), Adaptive Beta’s mean RTT was significantly lower than legacy’s ($d = 6.20$, $p < 0.0001$). The RTT Growth Guard showed a significantly *higher* mean RTT ($d = 5.80$, $p < 0.0001$), opposite its intended effect—plausibly because the freeze condition is evaluated against an RTT_{min} still settling during slow-start, causing stall-restart cycling. We flag this as a hypothesis, not a confirmed cause, requiring per-packet window-state logging to confirm (Section VI).




Fig. 7. Measured retransmissions for Legacy VenO, Adaptive Beta, and the RTT Growth Guard.

At 11 dB, Adaptive Beta’s mean goodput (10.14 Mbps) was statistically indistinguishable from legacy’s

(11.12 Mbps), and the Guard’s goodput was significantly lower (8.12 Mbps, $d = 2.59$, $p = 0.0041$). Neither amendment improved goodput at 11 dB, the most important unexplained outcome in this study: Adaptive Beta was designed specifically for this operating point, yet even a correctly-relaxed classifier could not recover goodput. We hypothesise that at 11 dB the link is intermittently exhausting its MAC retry budget, so even a correctly-classified “wireless” loss corresponds to a genuinely stalled link rather than a misclassification, and relaxing the classifier’s tolerance cannot recover goodput the MAC layer has not made available. This is proposed as a hypothesis for future kernel-level instrumentation, not a proven mechanism.

At the 5 dB floor (Fig. 7), both amendments showed significantly higher retransmissions than legacy (24.6 vs. 5.8, $d \approx 4.0$ – 4.5 , $p < 0.0001$). At 2 dB, Adaptive Beta alone sustained non-negligible goodput (0.30 Mbps vs. ≈ 0 ; $d = 21.89$), though values remain small.

Non-normality and variance heterogeneity mirrored the baseline dataset, so these comparisons are robust directional evidence rather than exact parametric guarantees.

VI. CONCLUSIONS AND FUTURE WORK

This work presented a fully empirical, VNA-calibrated evaluation of four TCP variants over a physically degraded Wi-Fi link, and used the resulting diagnosis of VenO’s static threshold to derive, implement, and empirically re-test two kernel-module amendments. Baseline results confirm Vegas’ latency advantage and Westwood’s goodput resilience near the transitional SNR boundary. Validation results are mixed: Adaptive Beta produced a significant RTT reduction at the clean-channel ceiling and uniquely sustained goodput at extreme exhaustion, but neither amendment improved goodput at 11 dB, and both showed higher retransmissions than legacy at 5 dB—a partial validation, not a net improvement.

Several scope limitations bound these conclusions. This study is restricted to a single-flow IBSS topology on legacy 802.11g, deliberately excluding AP scheduling and OFDMA as confounds; it therefore does not model contention-driven packet loss from random-access collisions among multiple stations sharing the ISM band, a principal loss mechanism in real deployments that is likely to interact differently with both classifiers than the single-link AWGN impairment studied here. The two amendments were also only tested independently; their interaction when active together was outside the scope

of this work and is proposed as future work below. Whether the β -scaling failure is specific to 802.11g MAC retry timing or manifests differently under 802.11n/ac/ax HARQ and OFDMA dynamics is likewise unaddressed, since the retry-latency assumptions underlying $k = 2.0$ (Section V-B) are specific to 802.11g. Finally, the Guard's unexpected 35 dB RTT increase remains open, plausibly linked to RTT_{min} drift during slow-start, and requires per-packet window-state logging to confirm or refute.

Future work should therefore: (i) empirically test the combined Adaptive Beta and RTT Growth Guard variant under a balanced statistical design; (ii) instrument RTT_{min} stability and the Guard's freeze/restart transitions at the kernel level; (iii) reintroduce multi-flow contention and random-access collisions as an experimental variable; and (iv) evaluate the amendment's constants against 802.11n/ac/ax MAC timing.

ACKNOWLEDGMENT

The authors thank the Faculty of Science, University of Malta, and C. Sammut, F. Cilia, and J. Bonello for technical assistance.

REFERENCES

- [1] M. Allman, V. Paxson, and E. Blanton, "TCP Congestion Control," RFC 5681, 2009.
- [2] S. Ha, I. Rhee, and L. Xu, "CUBIC: A new TCP-friendly high-speed TCP variant," ACM SIGOPS Oper. Syst. Rev., vol. 42, no. 5, pp. 64–74, 2008.
- [3] Y. Tian, K. Xu, and N. Ansari, "TCP in wireless environments: Problems and solutions," IEEE Commun. Mag., vol. 43, no. 3, pp. S27–S32, 2005.
- [4] G. Xylomenos, G. C. Polyzos, P. Mahonen, and M. Saura, "TCP performance issues over wireless links," IEEE Commun. Mag., vol. 39, no. 4, pp. 52–58, 2001.
- [5] A. Taha, A. E. Aziz, and M. Dessouky, "Comparative performance evaluation of TCP variants in WiMAX (and WLANs) network configurations," J. Comput. Netw. Commun., 2012.
- [6] S. Aoyagi, K. Yamamoto, and T. Hasegawa, "An accurate platform for investigating TCP performance in Wi-Fi networks," Future Internet, vol. 15, no. 7, p. 246, 2023.
- [7] X. Du et al., "Revisiting congestion control for Wi-Fi networks," in Proc. APNet, Sydney, Australia, 2024.
- [8] Q. Pang, S. C. Liew, C.-P. Fu, W. Wang, and V. O. K. Li, "Performance study of TCP Veno over WLAN and RED router," in Proc. IEEE WCNC, 2003.
- [9] S. Mascolo, C. Casetti, M. Gerla, M. Y. Sanadidi, and R. Wang, "TCP Westwood: Bandwidth estimation for enhanced transport over wireless links," Wireless Netw., vol. 8, pp. 467–479, 2001.
- [10] C. P. Fu and S. C. Liew, "TCP Veno: TCP enhancement for transmission over wireless access networks," IEEE/ACM Trans. Netw., vol. 11, no. 5, pp. 760–769, 2003.
- [11] L. Brakmo and L. Peterson, "TCP Vegas: End to end congestion avoidance on a global internet," IEEE J. Sel. Areas Commun., vol. 13, pp. 1465–1480, 1994.
- [12] IEEE Std 802.11g-2003, "Part 11: Wireless LAN MAC and PHY specifications: Further Higher Data Rate Extension in the 2.4 GHz Band," IEEE, 2003.
- [13] L. Friggieri Cordina, "Empirical Evaluation of TCP Variants over WiFi: Reproducibility Repository," GitHub, 2026. [Online]. Available: <https://github.com/LukeFriggieri/Empirical-Evaluation-of-TCP-Variants-over-WiFi>