

Title	The benefits of deceit: a malicious client in a 5G cellular network
Authors	Quinlan, Jason J.;Roedig, Utz
Publication date	2019-07
Original Citation	Quinlan, J. J. and Roedig, U. (2019) 'The benefits of deceit: a malicious client in a 5G cellular network', IEEE International Symposium on Local and Metropolitan Area Networks (LANMAN), Paris, France, 1-3 July, pp. 1-2. https://doi.org/10.1109/LANMAN.2019.8847162
Type of publication	Conference item
Link to publisher's version	10.1109/LANMAN.2019.8847162
Rights	© 2019, IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.
Download date	2026-08-31 05:05:21
Item downloaded from	https://hdl.handle.net/10468/8186



UCC

University College Cork, Ireland
 Coláiste na hOllscoile Corcaigh

The benefits of Deceit: a Malicious client in a 5G Cellular Network

Jason J. Quinlan and Utz Roedig

Computer Science & Information Technology, University College Cork, Ireland.

Email: [j.quinlan, u.roedig]@cs.ucc.ie

Abstract—As we advance towards Smart Cities, autonomous vehicles and the avalanche of IoT devices proposed for the future, we need to give careful consideration to how easily compromised nodes/devices can impact network state. Current proposals for autonomous smart devices typically use cellular networks as the backhaul or final hop. These devices will leverage existing trust-based client-side channel metrics, such as channel quality indicator (CQI), when the base-station determines scheduling decisions.

In this short paper, we investigate the scheduling impact of a malicious device when it changes its channel metrics, so as to improve its download rate or even to negate the download rate of others. We utilise real-time 4K ultra-high definition video delivery as an example of high throughput demand application and compare the delivery rates of multiple devices in an open-source 5G simulated NS-3 network. Our results illustrate that when a malicious client deceives the scheduler, the other clients in the network have a noticeable decrease in both viewable quality and underlying delivery rate (25% decrease in the average video quality across the non malicious clients).

I. INTRODUCTION

In the modern age, most network traffic originates from, interacts with or is destined for cellular devices. As we advance towards the billions of autonomous smart nodes, such as vehicular and IoT, destined for deployment over the coming years, cellular networks will become the first/last hop for practically all network traffic. The impact of how these devices will interact with network services, such as web traffic, video streaming, social media and messaging is forcing ISPs (cellular operators) to continue to improve their core network infrastructure and invest in new cellular technologies. The next iteration of the cellular technology is known as 5G, and promises high delivery rates, coupled with extremely low latency and overall improved reliability.

While network traffic can have inherent security at its core (*https*, *VPN*, *encryption*), the final cellular hop has a number of inherent trust-based metrics upon which traffic scheduling is determined. In current cellular networks, the download rate for a given client (allocated resource blocks) is typically determined by the base station using metrics from the device, such as the channel quality indicator (CQI). This mandates a level of trust between the device (the user) and the base station (ISP). As most current devices are not open-source hardware/software, modifications to the core functions of the device is limited, such that these trust-based metrics are inherently difficult to modify. With the introduction of open-sourced communication systems, such as software-defined

radio, as well as the multitude of IoT nodes being deployed, the rate of compromised or malicious devices in the network, with access to core functionality, will only increase.

In this short paper, we investigate the relationship between CQI value and allocated throughput, when a device maliciously change their channel metrics. We assume that an attacker cannot interfere with base stations or the network back end infrastructure and only has the ability to modify the behaviour of an end device. Specifically, the attacker can only modify performance data reported to the base station. The possible goals of the attacker include 1) Gain a better allocation of bandwidth than would be assigned normally, 2) Prevent other users from obtaining their fair bandwidth allocation and 3) Prevent other users from using services at all (push bandwidth allocation modification towards DoS).

II. BACKGROUND AND OVERVIEW

A low rate of malicious devices in a network can greatly impact on the Quality of Service (*QoS*) of all other devices within the range of the same base station. For 3G networks, it has been shown that if just 10% of the clients have malicious intent, a 2-second end-to-end inter-packet transmission delay for VOIP can be seen [1]. While sophisticated attacks can capture up to 77% of the allocated time slots [2]. For 4G, the outcome of the attacks is considerably worse, as just two malicious users can cause denial of service for all other users of the same cell requesting TCP-based applications [3].

These startling results, motivate us to investigate the impact of these malicious clients in a 5G network. We utilise an open-source simulated cellular testbed, namely the *DISGUISE* [4] framework for 5G evaluation. *DISGUISE* is a fully modifiable and extendable NS-3 based tool and offers a real-time hybrid physical and simulated infrastructure to evaluate high throughput applications. *DISGUISE* extends the NS-3 mmwave [5] module which offers a full-stack 5G infrastructure. Build and usage details for *DISGUISE*, are available on the website ¹.

In our evaluation, we utilise adaptive video delivery [6] as a means of illustrating how changing the download scheduling allocation can impact on the view-able quality perceived by a user. Adaptive video is designed so that for each streaming video clip, there exists a range of quality rates (average video encoding rate in Kbps). Each video quality rate is further broken into segments, where each segment permits

¹<http://www.cs.ucc.ie/misl/research/software/di5guise>

TABLE I
COLUMN NOTATION USED IN TABLE II AND TABLE III

Header	Description
Avg_quality	average view-able quality per segment in Kbps
Avg_del_rate	average download rate per segment in Kbps
Avg_del_time	average delivery time per segment in seconds
Num_stalls	number of stalls
Stall_Dur	total stall duration in seconds

a piece of the video to be downloaded (typically in the duration of seconds of content). Each rate offers options such as a different resolution, frame per second, and thus produce different perceived quality levels. This permits the client to select the correct quality rate segment which best reflects the cellular throughput rate seen at the device at that moment in time, so as to maximise perceived quality. In our simulated networks, both with and without malicious devices, we compare clients average download rate (Kbps), delivery time, the average view-able quality at the device (Kbps), and also highlight issues that can occur when the network changes drastically, such as video re-buffering and stalls. Investigating issues such as decoding the underlying modulation and coding scheme, spectral efficiency, and associated packet loss rates, is left for future work.

III. EXPERIMENTAL RESULTS

In our NS-3 5G mmwave evaluation, four clients range between 25m and 35m in distance from a single base station. We utilise a rural macro-channel scenario (*RMa*), with a 3GPP propagation loss model (*3GPPprop*), and a Transmission Time Interval (*TTI*) mac scheduler. For all evaluation scenarios, we stream five minutes of 4K adaptive video content to a number of video clients, who will adapt the quality of their streaming video clip depending on their scheduled throughput rate. For the evaluation with a malicious client, we maximise the CQI value of Client one. The notation used in the column headers of Table II and Table III is shown in Table I.

Table II illustrates the client metrics when there are no malicious clients in the network. We can see that the average quality rate per segment across all four clients is approximately 11,388 Kbps, which is reflected in at least two of the clients (denoting a fair allocation across all user space). We see similar equality over delivery rate and delivery time for all four clients. Delivery time is reflective of average video quality as higher quality takes longer to download. There is noticeable stalls in three of the clients, but this can occur when the client makes the wrong decision for the next video segment and spends too long waiting for the segment to arrive (typically occurs when the network throughput drops and the client is stuck downloading a segment with a large transmission cost).

In Table III reflects the client metrics when client 1 is malicious and maximise the CQI value sent to the scheduler. We not only see a much lower average quality rate per segment across all four clients, approximately 9,290 Kbps, but also a marked decrease in the overall average quality rate for the non malicious clients. These lower average values are also noted in the delivery rate and delivery time of Clients two to four. One

TABLE II
5G CLIENT METRICS WITH NO MALICIOUS CLIENTS

Client	Avg_quality	Avg_del_rate	Avg_del_time	Num_stalls	Stall_Dur
1	11,508	20,091	2.435	0	0
2	8,934	16,589	2.147	1	3.337
3	13,477	22,141	2.738	2	5.22
4	11,636	20,618	2.568	3	3.395

positive item is that two of the client no longer stall during playback, but Client 3 has seen a 100% increase in both stall events and overall duration.

TABLE III
5G CLIENT METRICS WITH CLIENT 1 AS THE MALICIOUS CLIENT

Client	Avg_quality	Avg_del_rate	Avg_del_time	Num_stalls	Stall_Dur
1	11,997	20,221	2.706	0	0
2	7,246	14,823	2.064	0	0
3	9,865	17,318	2.433	4	10.422
4	8,144	16,030	2.221	0	0

IV. DISCUSSION AND CONCLUSION

Discussion: It is interesting to note that while client 1 is maximising the CQI value sent to the scheduler, there is only a 4% increase in average video quality. This increase does not reflect the higher resource block allocation by the scheduler, and the marked decrease in video quality of the other client. Further work is required to extract the relevant features from the 5G scheduler, so as to better compare this relationship.

Conclusion: In this paper we investigated the impact on 4K adaptive video quality in a multi-user 5G network, when a malicious client utilises existing trust-based channel metrics to deceive the cellular scheduler so as to maximise its download bandwidth allocation.

ACKNOWLEDGMENT

The authors acknowledge the support of Science Foundation Ireland (SFI) under Research Grant 13/IA/1892.

REFERENCES

- [1] K. Pelechrinis, P. Krishnamurthy, and C. Gkantsidis, "Trustworthy operations in cellular networks: The case of pf scheduler," *IEEE Transactions on Parallel and Distributed Systems*, vol. 25, no. 2, pp. 292–300, Feb 2014.
- [2] R. Racic, D. Ma, H. Chen, and X. Liu, "Exploiting and defending opportunistic scheduling in cellular data networks," *IEEE Transactions on Mobile Computing*, vol. 9, no. 5, pp. 609–620, May 2010.
- [3] R. Bassil, I. H. Elhajj, A. Chehab, and A. Kayssi, "A resource reservation attack against lte networks," in *2013 Third International Conference on Communications and Information Technology (ICCIT)*, June 2013, pp. 262–268.
- [4] J. J. Quinlan, K. K. Ramakrishnan, and C. J. Sreenan, "DI5GUISE: A highly Dynamic Framework for Real-Time Simulated 5G Evaluation," in *Proc. of 25th IEEE International Symposium on Local and Metropolitan Area Networks (LANMAN 2019)*, 2019.
- [5] M. Mezzavilla, M. Zhang, M. Polese, R. Ford, S. Dutta, S. Rangan, and M. Zorzi, "End-to-End Simulation of 5G mmWave Networks," *IEEE Communications Surveys Tutorials*, vol. 20, no. 3, pp. 2237–2263, thirdquarter 2018.
- [6] J. J. Quinlan and C. J. Sreenan, "Multi-profile Ultra High Definition (UHD) AVC and HEVC 4K DASH Datasets," in *Proceedings of the 9th ACM Multimedia Systems Conference, ser. MMSys '18*. New York, NY, USA: ACM, 2018, pp. 375–380. [Online]. Available: <http://doi.acm.org/10.1145/3204949.3208130>