

Poster: Redundancy-Aided Vehicular Networking

HyunJong Lee

Jason Flinn

University of Michigan

Vehicular applications are increasingly connected to cloud services. For example, route planning, gas price applications, and Siri-like personal assistants all respond to user queries based in part on cloud processing. Network communication thus is often a substantial component of user-perceived latency in vehicular applications.

Current vehicular computing platforms typically connect to the cloud using a single cellular network provider. Network conditions can change rapidly as a vehicle moves due to geographical variation in coverage, radio shadows, and differing traffic density. Such variation is often exacerbated by connection re-establishment after an interface has entered a sleep state.

Thus, vehicular applications can often appear unresponsive due to high wireless network latency. Even worse, the responsiveness is unpredictable; high tail latency makes some user interactions take longer, even when most interactions complete in an acceptable amount of time. This unpredictability is especially worrisome in a vehicular environment, in which occasional unexpected performance anomalies distract the driver of the vehicle.

Smartphones mitigate poor network performance by switching among multiple networks (e.g., WiFi and cellular). Such a strategy makes sense when performance is predictable. However, it provides little help for the unexpected tail latencies common in vehicular environments.

In our work, we are mitigating tail latency for vehicular applications by redundantly sending data over multiple wireless networks for user-facing applications. Rather than estimate network latency and bandwidth as scalar values, we calculate the expected distribution of these values for each wireless networks. If the error bounds for the expected time to transmit small amounts of data over at least two different wireless networks overlap with each other, our vehicular platform transmits the data over both networks. At the cost of redundant transmission of data, this strategy improves the average user-perceived delay in vehicle-cloud applications. More importantly, it substantially reduces tail latency.

The opportunities for employing redundant transmission are considerable in vehicular computing. Not only can the vehicular platform employ built-in cellular and WiFi inter-

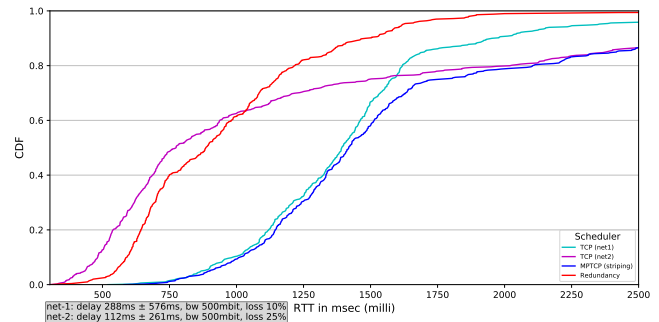


Figure 1: 10K data transfer by different schedulers

faces, but the platform can also communicate with multiple cellular providers by tethering with smartphones carried by passengers within the vehicle.

We argue that the downsides of employing redundancy are surprisingly small. The user-facing applications for which we are employing redundancy transmit small amounts of data. Indeed, an application transmitting bulk data would benefit little from redundancy, and our vehicular platform would not employ redundancy for such a purpose. Thus, the total amount of bytes sent redundantly is small. Energy considerations are also minimal since vehicular platforms are typically powered by the automobile engine.

Our prototype system for *Redundancy-Aided Vehicular Networking* (RAVEN) reduces tail latency for vehicular applications via redundant network transmission. RAVEN is implemented on top of MPTCP, a recent extension of TCP to support multipath networks. MPTCP binds a single application TCP connection to multiple TCP subflows to stripe data over multiple physical networks. RAVEN implements a new MPTCP scheduler that selectively transmits small data redundantly over these subflows. We have developed a statistical model that provides latency estimates for each physical network within a configurable confidence level. When redundancy has a substantial confidence to reduce tail latency, RAVEN transmits the data redundantly over multiple subflows to improve responsiveness of user-facing applications.

We are currently evaluating RAVEN in many wireless network scenarios. For instance, consider the common situation where two wireless networks are experiencing unexpected packet loss. As shown in Fig. 1, RAVEN outperforms both MPTCP and a strategy that selects either wireless network for exclusive transmission. The improvement in performance is most noticeable at the 99% percentile in the CDF, where redundancy substantially reduces tail network latency compared to generic MPTCP and TCP over either network.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author(s). Copyright is held by the owner/author(s).

MobiSys '17, June 19–23, 2017, Niagara Falls, NY, USA.

ACM ISBN 978-1-4503-4928-4/17/06.

DOI: <http://dx.doi.org/10.1145/3081333.3089317>