Simulation Modelling Practice and Theory 000 (2016) 1–12

Contents lists available at ScienceDirect

Simulation Modelling Practice and Theory

journal homepage: www.elsevier.com/locate/simpat

Computation offloading of a vehicle's continuous intrusion detection workload for energy efficiency and performance

George Loukas*, Yongpil Yoon, Georgia Sakellari, Tuan Vuong, Ryan Heartfield

Computing and Information Systems, University of Greenwich, UK

ARTICLE INFO

Article history:
Available online xxx

Keywords:
Computation offloading
Intrusion detection
Energy efficiency
Detection latency
Cyber-physical systems
Vehicular security

ABSTRACT

Computation offloading has been used and studied extensively in relation to mobile devices. That is because their relatively limited processing power and reliance on a battery render the concept of offloading any processing/energy-hungry tasks to a remote server, cloudlet or cloud infrastructure particularly attractive. However, the mobile device's tasks that are typically offloaded are not time-critical and tend to be one-off. We argue that the concept can be practical also for continuous tasks run on more powerful cyber-physical systems where timeliness is a priority. As case study, we use the process of real-time intrusion detection on a robotic vehicle. Typically, such detection would employ lightweight statistical learning techniques that can run onboard the vehicle without severely affecting its energy consumption. We show that by offloading this task to a remote server, we can utilse approaches of much greater complexity and detection strength based on deep learning. We show both mathematically and experimentally that this allows not only greater detection accuracy, but also significant energy savings, which improve the operational autonomy of the vehicle. In addition, the overall detection latency is reduced in most of our experiments. This can be very important for vehicles and other cyber-physical systems where cyber attacks can directly affect physical safety. In fact, in some cases, the reduction in detection latency thanks to offloading is not only beneficial but necessary. An example is when detection latency onboard the vehicle would be higher than the detection period, and as a result a detection run cannot complete before the next one is scheduled, increasingly delaying consecutive detection decisions. Offloading to a remote server is an effective and energy-efficient solution to this problem too.

1. Introduction

Offloading computation tasks from a user's device to a remote server, cloudlet or cloud [1] can have multiple benefits. It can allow the utilisation of more powerful and more flexible computing resources and can provide an on-demand service, while also dramatically reducing the energy cost on the user's device. For these reasons, it has evolved into common practice for mobile devices [2]. We argue that for largely the same reasons, the concept of computational offloading can be extremely useful for demanding, real-time and continuous tasks required by more powerful yet still resource-constrained and time-critical cyber-physical systems, such as vehicles. Yet, the approach remains largely unexplored. In the context of smart cities, it is anticipated that in the near future, the majority of vehicles on the road will benefit from various forms of constant

* Corresponding author.
E-mail address: g.loukas@gre.ac.uk (G. Loukas).

http://dx.doi.org/10.1016/j.simpat.2016.08.005
1569-190X/

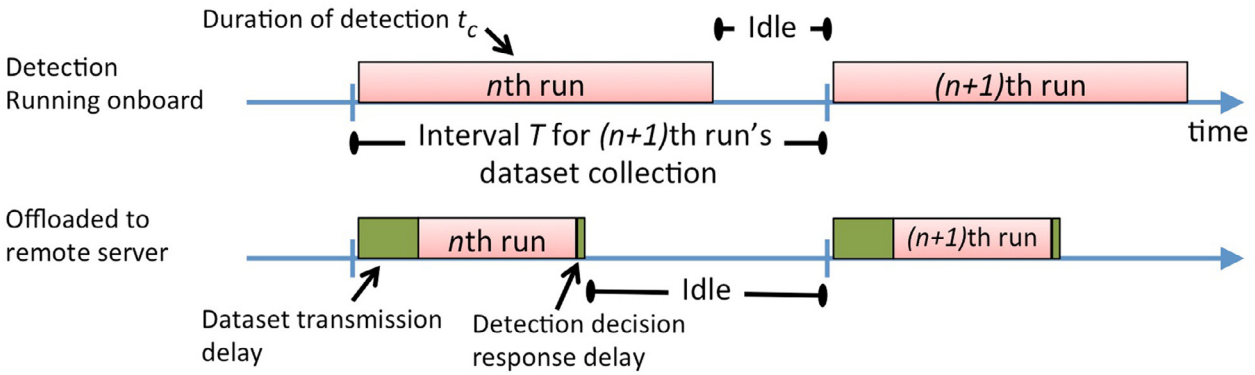

Fig. 1. To ensure continuous coverage, the periodic detection interval needs to be roughly equal to the periodic data collection interval. In the offloading case (bottom), the delay of transmitting the dataset over the network needs to be taken into account.

connectivity with smart city infrastructures, as well as with each other. Already a wide range of vehicle-to-infrastructure (V2I) and vehicle-to-vehicle (V2V) technologies for safety and comfort have been developed, with some already deployed operationally [3]. However, the dependence on V2I and V2V communications generates several new cyber threats to vehicles operating in smart cities. Protection against a wide range of evolving smart city cyber threats may be impractical if relying only on the vehicle's onboard processing systems, because of the adverse impact it would have on energy efficiency and latency. This makes computation offloading highly relevant in this context.

As proof of concept, we focus on the provision of high-performance intrusion detection for a robotic vehicle with limited processing resources and a requirement for low detection latency. A key challenge here is that to be meaningful the task needs to be performed continuously, which can be highly impractical for a battery-powered robotic vehicle. At the same time, it needs to provide a detection decision in time before it receives the next data to analyse. This is also a major challenge for advanced detection mechanisms that exhibit high detection latency and become impractical if they detect an attack after it has physically damaged the vehicle, e.g. through a command injection attack that disables or selectively engages the brakes on one side [4,5]. Here, we develop a prototype system to demonstrate that it is possible to employ offloading in this case to reduce not only the energy cost but also the overall time taken to complete the task, even when the additional networking and processing overheads are taken into account. In several cases experimented with, reducing detection latency is not only useful, but necessary to ensure that detection can still be technically practical for a real-time system.

The practicality of computation offloading relies generally on two factors: performance and energy cost. For the particular task, performance relates to overall detection latency. In the onboard detection case, this effectively corresponds to the time it takes to complete the computation for detection (Fig. 1: top). In the offloading case, it includes the time it takes to send the data to the server over the network, the server to complete the computation and the time to receive the result back from the server (Fig. 1: bottom). However, the usual assumption in modelling offloading is that the latter delay is insignificant in comparison to the first two due to the difference in size between the data transmitted for offloading and the response. In our case, the response is simply a binary value (1 being an attack and 0 a non-attack).

As energy is power times time, the energy cost due to detection depends on the additional power consumed on the vehicle when running the detection computation onboard, as well as the time taken to complete it. In the offloading case, it also includes the energy cost of transmitting the data to the remote server, as transmission power consumption tends to be higher than idle operation power consumption. As with time, the energy cost of receiving the response can also be disregarded.

The key contributions of this work are: (i) A proof of concept prototype for computation offloading to provide a vehicle with access to high-end machine learning algorithms, with a case-study in deep learning for intrusion detection, (ii) a mathematical model of the difference in energy costs between onboard and offloaded computation for continuous periodic tasks, which is validated experimentally, and (iii) an experimental evaluation of the energy consumption and detection latency for a robotic vehicle by offloading deep learning-based detection mechanisms of low, moderate and high complexity[1].

2. Related work

Computation offloading has been thoroughly studied for mobile devices, in terms of both performance and energy efficiency, but to a very limited extent for IoT and cyber-physical systems, especially with regards to energy efficiency. The following is a brief overview of related work in these three areas.

[1] Please note that we are using the word 'complexity' rather loosely, meaning that the time a specific deep-learning detection mechanism takes to complete increases as we add more features and more hidden neurons (hence moving from a lower to a higher complexity deep learning model)

Please cite this article as: G. Loukas et al., Computation offloading of a vehicle's continuous intrusion detection workload for energy efficiency and performance, Simulation Modelling Practice and Theory (2016), http://dx.doi.org/10.1016/j.simpat.2016.08.005