Demo: Distributed Real-Time Generative 3D Hand Tracking using Edge GPGPU Acceleration

Ammar Qammaz
FORTH and University of Crete
Heraklion, Crete, Greece
ammarkov@ics.forth.gr

Sokol Kosta
Aalborg University Copenhagen
Copenhagen, Denmark
sok@es.aau.dk

Nikolaos Kyriazis
Insitute of Computer Science, FORTH
Heraklion, Crete, Greece
nkyriazis@gmail.com

Antonis Argyros
FORTH and University of Crete
Heraklion, Crete, Greece
argyros@ics.forth.gr

ABSTRACT

This work demonstrates a real-time 3D hand tracking application that runs via computation offloading. The proposed framework enables the application to run on low-end mobile devices such as laptops and tablets, despite the fact that they lack the sufficient hardware to perform the required computations locally. The network connection takes the place of a GPGPU accelerator and sharing resources with a larger workstation becomes the acceleration mechanism. The unique properties of a generative optimizer are examined and constitute a challenging use-case, since the requirement for real-time performance makes it very latency-sensitive.

3D REAL-TIME HAND TRACKING

3D hand tracking is a complex computer vision problem with 27 degrees of freedom. In order to perform hand tracking in real-time, devices need high-end GPUs, given that the algorithms used to solve this problem are implemented using CUDA. In this demo, we make use of edge computation offloading to remotely execute the CUDA operations from thin clients to more powerful servers with GPU access. We deploy an Edge Architecture and use the RAPID framework [1], which automatically monitors and offloads the heavy workloads depending on the overall conditions of execution. This also enables thin clients to execute code they would otherwise not be able to, due to lack of sufficient hardware.

The demonstration focuses on performing Real-Time 3D Hand Tracking [2] on a low-end laptop that receives input through a USB ASUS Xtion RGBD sensor. The heavy GPGPU processing is then automatically offloaded via the RAPID framework to a faster machine that features a GPGPU and acts as a server that executes all intensive calculations. Figure 1 depicts a snapshot from demo execution[1, 2]. Two different laptops are involved in the demo. The high-end laptop features a state-of-the-art NVIDIA GTX1080M graphics card, while the low-end one has a dated NVIDIA GTX670M card from 2012, which is no longer supported by NVIDIA (due to compute capability 2.1 being dropped in CUDA 9.0) and used as the thin client. The machines are connected via Gigabit Ethernet and WiFi and different offloading strategies can be requested to showcase the performance achieved in the various network topologies. The Hand Tracker output is augmented on the RGB feed of the camera and visualization is displayed in the screen of the laptop to provide feedback to the user. Using this method: *i)* Devices that do not have GPGPU hardware can leverage resources from high-end devices through the network. *ii)* The framerate achieved depends on the connection quality with the best results achieved when using a low-latency Gigabit Ethernet connection. *iii)* The code only requires a few modifications to the original application to become remoteable via the provided RAPID API [1].

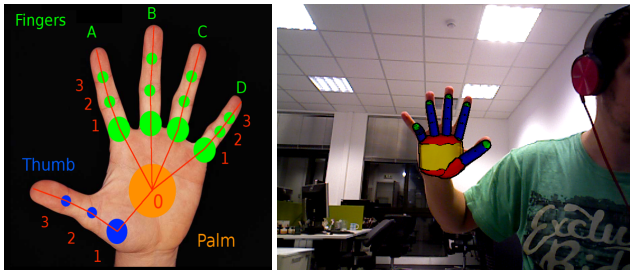

Figure 1: 3D Hand Tracking demo overview.

Acknowledgement: The work was partially supported by the EU Horizon 2020 programs through the RAPID (H2020-ICT-644312) and Co4Robots (H2020-ICT-2016-1-73186) projects.

REFERENCES

[1] R. Montella, S. Kosta, D. Oro, J. Vera, C. Fernández, C. Palmieri, D. Di Luccio, G. Giunta, M. Lapegna, and G. Laccetti, "Accelerating linux and android applications on low-power devices through remote gpgpu offloading," *Concurrency and Computation: Practice and Experience*, vol. 29, no. 24, 2017.

[2] I. Oikonomidis, N. Kyriazis, and A. A. Argyros, "Efficient model-based 3d tracking of hand articulations using kinect," in *British Machine Vision Conference (BMVC 2011)*, vol. 1, no. 2. Dundee, UK: BMVA, 2011, pp. 1–11.

[3] A. Qammaz, S. Kosta, N. Kyriazis, and A. Argyros, "On the Feasibility of Real-Time 3D Hand Tracking using Edge GPGPU Acceleration," *ArXiv e-prints*, Apr. 2018.

[1]Video: https://youtu.be/CBAjA6OUFnc
[2]Full paper: https://arxiv.org/abs/1804.11256 [3].

MobiSys '18, June 10–15, 2018, Munich, Germany

ACM ISBN 978-1-4503-5720-3/18/06.
https://doi.org/10.1145/3210240.3211112