



Aalborg Universitet

AALBORG UNIVERSITY
DENMARK

On the Feasibility of Real-Time 3D Hand Tracking using Edge GPGPU Acceleration

Qammaz, A.; Kosta, S.; Kyriazis, N.; Argyros, A.

Publication date:
2018

Document Version
Publisher's PDF, also known as Version of record

[Link to publication from Aalborg University](#)

Citation for published version (APA):

Qammaz, A., Kosta, S., Kyriazis, N., & Argyros, A. (2018). *On the Feasibility of Real-Time 3D Hand Tracking using Edge GPGPU Acceleration*. arXiv.org. <https://arxiv.org/abs/1804.11256>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- ? Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- ? You may not further distribute the material or use it for any profit-making activity or commercial gain
- ? You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

On the Feasibility of Real-Time 3D Hand Tracking using Edge GPGPU Acceleration

Ammar Qammar
FORTH and University of Crete
Heraklion, Crete, Greece
ammarmov@ics.forth.gr

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

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

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

ABSTRACT

This paper presents the case study of a non-intrusive porting of a monolithic C++ library for real-time 3D hand tracking, to the domain of edge-based computation. Towards a proof of concept, the case study considers a pair of workstations, a computationally powerful and a computationally weak one. By wrapping the C++ library in Java container and by capitalizing on a Java-based offloading infrastructure that supports both CPU and GPGPU computations, we are able to establish automatically the required server-client workflow that best addresses the resource allocation problem in the effort to execute from the weak workstation. As a result, the weak workstation can perform well at the task, despite lacking the sufficient hardware to do the required computations locally. This is achieved by offloading computations which rely on GPGPU, to the powerful workstation, across the network that connects them. We show the edge-based computation challenges associated with the information flow of the ported algorithm, demonstrate how we cope with them, and identify what needs to be improved for achieving even better performance.

1 INTRODUCTION

Mobile devices have become ubiquitous in western societies, networking technologies are constantly improving, and cloud computing services offer an abundance of resources for applications that support them. On the other hand, factors such as cost of purchase, power consumption, battery life, as well as the large number of legacy mobile devices still being used by consumers, make a large portion of the mobile computer market underpowered relative to the top-tier devices commercially available. Of course, high-end devices such as gaming laptops do exist and offer amazing performance, but their cost is prohibitive for the average consumer and they mainly cater to a smaller gamer niche. This leaves most devices equipped with cheap hardware alternatives such as on-board graphics and slower and more power conserving CPUs. This, in turn, translates to less computing resources for applications albeit in more affordable computers. Another trend, which is one of the motivations behind this work, is that computer users increasingly view the world from their devices with a surprising number of new HCI devices recently entering the market, mostly in the form of virtual assistants such as Amazon Echo, Google Home, Siri, and Kortana, that currently lack vision and rely on voice commands and



Figure 1: Indicative hand tracker output overlaid over the RGB feed of an RGBD camera source.

only provide audio output using a network back-end. In a world that is becoming increasingly connected, we believe that humans will rely more and more on computer vision techniques to facilitate human-computer interaction. These trends hint that this technology brings great future potential for applications that can cater to the needs of users by leveraging the tools currently available. As an example, a robust, vision-based hand tracker module enables a computer system to recover the 3D position, orientation and full articulation of a human hand unobtrusively, from visual input. This is a stepping stone for building additional components such as gesture recognition. If this functionality becomes accessible to low-end devices, a variety of tasks (as simple as controlling a media player or as complex as household robotics and managing smart-homes by using hand signs), can be performed via simple, intuitive gestures.

In this work, we report on the experience and challenges of extending a powerful generative real-time 3D Hand Tracker to make it exploit Edge Offloading for achieving satisfactory results even when running on low-power devices.

The rest of the paper is organized as follows: Section 2 presents an overview of the most relevant works in the field; Section 3 describes the details of the GPU-based implementation of the hand tracking system and its adaption for low-power devices thanks to Edge Offloading; Section 4 presents the results of the experiments performed in different configuration setups; and finally, Section 5 concludes the paper with final remarks on the lessons learned by running a real-time application on the edge via computation offloading.

2 RELATED WORK

We review the research works that are most relevant to real-time hand tracking and network acceleration/computation offloading.

2.1 Hand Tracking

Real-Time 3D Hand Tracking is an open research problem in Computer Vision. Several approaches have been proposed towards achieving an effective and efficient solution. A very recent overview of such research efforts and state of the art hand tracking solutions is presented in [1]. Methods can be roughly categorized based on the camera sensors they use and on the adopted hand pose estimation approach. These approaches are broadly classified into discriminative, generative, and hybrid. Recent discriminative works use Convolutional Neural Networks (CNNs) with off-the-shelf RGB cameras but as seen in [1], 3D volumetric representations outperform 2D CNNs and 2D CNNs are not directly suitable for 3D hand pose estimation [2]. We thus focus on real-time methods that use RGBD/depth sensors. Discriminative methods such as [3] use statistical modeling, other works [2, 4] use 3D CNNs, while other solutions use cascades [5] or forest based classifiers [6]. Generative approaches to the problem involve the use of particle filters [7] or genetic algorithms [8] to perform regression and solve the problem. Hybrid solutions with both discriminative and generative components have also been proposed [9, 10]. It is important to note that almost all state of the art methods utilize GPGPU acceleration in order to achieve high framerates.

2.2 Automatic Network Acceleration

Computation offloading is a simple but powerful concept, which has been used successfully to help thin client devices execute resource and energy-hungry applications [11]. The first proposed solutions, such as Cloudlets [12], MAUI [13], CloneCloud [14] and ThinkAir [15], among others, have shown that it is beneficial for low-powered devices to offload the heavy tasks to remote powerful entities for remote execution, under the right circumstances. Indeed, these works show that the computation offloading works well when at least one or more of the following conditions are satisfied: *i*) the task to be offloaded is computationally heavy, *ii*) the amount of data to be transferred is limited, and *iii*) the network between the thin client and the powerful remote entity is good. For this reason, the applications considered by these works can be classified as CPU-hungry, delay-tolerant (up to a certain level, meaning they are not real-time), and not data-intensive.

With the recent progress in the distribution of networking, storage, and computing resources, i.e. increasing number of data centers around the world, and with the advent of the 5G network, more advanced offloading frameworks have been proposed, which deal with applications that present also challenging features, such as real-time or data-intensive. Flores et al. [16] present a framework for computation offloading from Internet of Things (IoT) devices, with a solution for load-balancing on the cloud side when handling multiple requests for task offloading. Chatzopoulos et al. [17] propose OpenRP, a platform for avoiding selfish behavior in the context of Device-to-Device offloading, where mobile devices can help each-other execute tasks in a distributed manner.

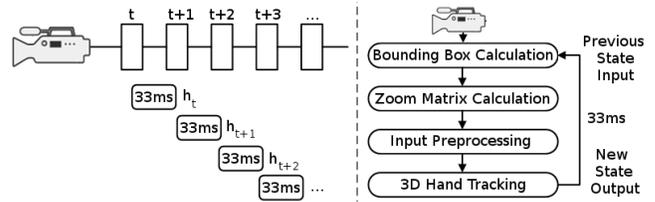


Figure 2: Left: A real-time framerate of 30fps means that every frame acquired by a camera has to be consumed/processed in less than 33 milliseconds. Right: 3D hand tracking consists of four discrete optimization steps which need to be executed in the 33ms time window. We use a JNI container that exposes these steps to the offloading framework and can choose to offload each of them separately (Multiple-Step experiments, Figure 5) or all of them in a single step (Single-Step experiments, Figure 5) in order to reduce intermediate data transfers and network latency.

Montella et al. [18] are the pioneers of open source solutions for GPU code offloading in low power devices. RAPID¹ is a European project that provides a framework for CPU and GPU computation offloading from low-powered devices to the edge of the network or to the cloud, depending on the context [19, 20]. In the current work, we make use of the RAPID tools to implement and test the GPU-based real-time hand tracking on thin clients.

3 3D HAND TRACKING ON THE EDGE

3.1 Generative 3D Hand Tracking

The hand tracker application is a processor of frames generated by a camera at a framerate of 30 frames per second (fps). It can be considered as a black box that receives a hand configuration h_t at time t (hand 3D position, orientation, articulation), along with a frame pair of RGB c^o and depth d^o images at time $t + 1$ and computes an estimation h_{t+1} of the hand configuration for the input frame. Repeating the procedure for every acquired RGBD frame results in the tracking of the observed hand.

Hand model: Each hand is encoded as a vector h of 27 kinematic parameters. 3 parameters represent its 3D location, and 4 more parameters represent its 3D orientation using a right-hand Cartesian coordinate system (X, Y, Z) and using a quaternion representation to avoid gimbal locks. The rest 20 parameters represent bone angles that encode finger articulation.

Objective function: An RGBD frame is denoted as $o = (c^o, d^o)$, where c^o and d^o stand for the camera RGB and depth frames, respectively. Moreover, given an instance h of the hand model and the camera calibration parameters, we can render the hand model to the camera viewport, obtaining color and depth maps $r = (c^h, d^h)$ that are directly comparable to the observations. An objective function $E_D(h, o)$ can then be defined that quantifies the discrepancy between a hand configuration / model hypothesis h and the actual observations o . Estimating the hand pose at a certain frame amounts

¹<http://www.rapid-project.eu/>

to finding the model parameters h^* that minimize $E_D(h, o)$:

$$h^* \triangleq \arg \min_h E_D(h, o). \quad (1)$$

More specifically, the error term that quantifies the discrepancy between a hand configuration h and the available observations, o is defined as:

$$E_D(h, d^o) = \frac{1}{N_p} \sum_{p \in B} C(|d_p^h - d_p^o|, T). \quad (2)$$

In Eq.(2), E_D sums the absolute depth differences $|d_p^h - d_p^o|$ for all points p that belong to a bounding box B containing the hand. The clamping function $C(x, T)$ returns x if $x \leq T$ and T otherwise. This is used to make the error term robust and prevent outliers from affecting it too much. In our implementation we set $T = 30\text{cm}$.

Hand pose estimation in a certain frame can thus be thought of as an optimization/minimization problem that seeks for the best set of 27 parameter values that result in a hand pose hypothesis that best matches the available observations.

Particle Swarm Optimization (PSO): Particle Swarm Optimization (PSO) [21] is a stochastic method that performs optimization by iteratively improving a candidate solution with respect to an error term characterizing its quality (objective function). PSO maintains a population of candidate solutions, called particles, that have a position and a velocity in the search space. The movement of each particle is influenced by the best position this particle has ever visited up to the current iteration/generation, and simultaneously guided towards the globally best known position in the search space (i.e., the best of the best for all particles). This strategy moves the swarm toward the best solutions.

PSO does not require training and does not need to compute the gradient of the optimized objective function which, thus, does not need to be differentiable. PSO is ideal for parallel implementation on modern GPU architectures, permitting interactive framerates and a GPGPU implementation provides 100× speedup compared to a serial implementation.

PSO has been applied successfully to a number of vision problems such as object detection [22], head pose estimation [23], 3D hand tracking [8, 24], 3D tracking of hands in interaction with objects [25–27], estimating contact forces [28], as well as 3D human pose tracking [29–31]. Tracking more complex object configurations than a single hand increases the search space of PSO exponentially. As an example, tracking two hands results in a 54-D search space, and each subsequent rigid object adds 7 more dimensions (its position plus its orientation) [25, 27]. In turn, the increase of the dimensionality of the problem creates the need for more computational resources and a larger optimization budget (i.e., number of evaluations of candidate solutions). Offloading to powerful machines, with more computational resources than a thin client, is expected to help towards an effective solution to this problem. In

this work, we examine the basic 27 degree of freedom (dof) single hand-tracking problem [8], since it represents the first step towards addressing the challenges of real-time computation offloading in the edge for facilitating Human-Computer Interaction (HCI) and other real-life scenarios.

PSO for 3D hand tracking: For our 3D hand tracking problem, PSO minimizes the objective function of Eq.(1) over candidate solutions h . For each incoming frame, particles are initialized around the solution of the previous frame. The space around that solution is made large enough to include the current frame estimation.

The optimization internally happens in 4 consecutive steps, as depicted in Figure 2. In order to achieve a 30 fps rate, and process every frame in real-time, the system should perform every processing step in less than 33 ms. Otherwise, if the system takes more than 33 ms per step, the whole pipeline gets slowed down, because the hand configuration h_{t+1} cannot be estimated without considering the estimation h_t of the hand in the previous time step. This is depicted in segment A of Figure 3, where we can observe that for a hypothetical slower 150 ms processing loop time, the system must skip processing two consecutive frames for each received frame, since it does not have enough processing time for them. Dropping frames is bad not only for the user experience (as there is observable delay), but also for the quality of the tracking. During the time unaccounted for, the hand moves further away from the last tracked/estimated position. This means that the hand tracker has to compensate by considering hand configurations in a much wider area in the space of hand poses, which makes the problem much more difficult. Additionally, errors tend to accumulate and can quickly lead to track loss in case of fast motion and low frame rates.

This also means that we cannot simply transmit a video feed over the network and just collect tracking results at the rate of acquisition. Instead, our client needs to wait for a frame to be processed and then return back with the best explanation h in order to submit the next one and start a new search. This makes it a particularly hard case since we are aiming at real-time performance but any latencies experienced through the network medium are aggregated to the overall loop time of the client.

3.2 3D Hand Tracking on Thin Clients

We use the RAPID platform as the offloading mechanism for enabling real-time hand tracking on thin clients. RAPID is a framework for Java and Android low-power devices, which can be used for both CPU and GPGPU code offloading. As such, given that the Hand Tracker was originally written as a C++ library, we created a Java Native Interface (JNI) container for the library, as well as a Java front-end that allowed us to achieve automatic offloading using RAPID. RAPID’s use of the maven build system allow effortless packaging of all the required runtime-libs of the application across machines, and the Java front-end makes it very easy to control and script it on a high-level without exposing the developer to the complexity of tens of thousands lines of code underneath.

When the RAPID-enabled application starts for the first time, a registration process is triggered, which involves sending the bytecode of the application to the remote side of the RAPID system.

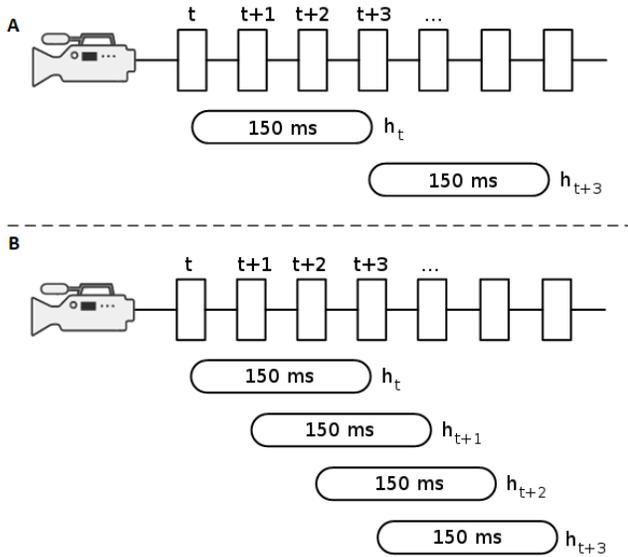


Figure 3: A: Algorithms with serial frame dependencies have to wait until a frame is processed before consuming the next. B: An offloading architecture can accommodate much better rates when not having to wait for a request to finish. Generative 3D Hand Tracking falls in category A.

Being a method-level offloading framework, whenever an offloadable Java method is called, RAPID decides at runtime if the method should be executed locally on the device or remotely on the powerful machine. If the offloaded Java method uses *native C/C++* functions, which is the case of our application, the RAPID remote counterpart will dynamically load the shared libraries (*.so* files) embedded in the application, which were transmitted during the registration phase, and will try to resolve the native function. Thanks to this feature, we were able to easily utilize the offloading framework straightforwardly. Moreover, if a native function contains also CUDA calls, which again is the case of our application, they are automatically executed on the GPU as if they were running normally without any offloading involved.

Thanks to offloading, we are not constrained by the resources of a single machine, which could potentially allow us to scale up as seen in Figure 3 on part B of the diagram. Unfortunately, as described in the previous sections, the nature of the optimization framework used for the hand tracker is not suited for this kind of parallel processing, given that each frame must first be processed and produce the output before the next one can be processed. Thus, we lose the benefits of concurrency and any potential benefit gained by assigning each incoming frame to a separate computing resource that can accommodate it, given that the system needs to wait for each step to be completed before handling the next step.

The source-code of the Java/JNI wrapper implementation is distributed as open source and can be found in GitHub [32]. The C++ Hand Tracker native library is also provided in binary form in the repository, allowing for the reproduction of the following experiments.

4 EXPERIMENTAL EVALUATION

4.1 Experiment Setup

In order to study the behavior of our RAPID-enabled hand tracker application, we use a testing environment with two different tiers of devices: a high-end server and a low-end laptop. The high-end server features a GeForce GTX 1080M GPU and an Intel Core i7 processor, while the laptop has an outdated GeForce 670M GPU and an early generation Intel Core i5.

Moreover, we connect the devices using two networks: a fast Gigabit Ethernet connection and a slower 802.11 Wi-Fi connection.

As mentioned above, to achieve a 30 fps tracking loop rate, which allows to process every frame received from the RGBD device, all the processing needs to be executed within 33 ms. Unfortunately, Wi-Fi connections are very prone to radio interference and typically introduce latency ranging from 10–60 ms, depending on the number of connected clients and network saturation. Moreover, the available bandwidth of a Wi-Fi connection is substantially lower than the one provided by a Gigabit Ethernet connection.

In order to have comparable results with the different setups, we pre-recorded a video depicting various challenging hand movements.

Having the same input stream to evaluate across all runs, we aim to identify the overhead introduced by the network connections, by the RAPID offloading framework, and by the Java’s JNI wrapper and the Java Virtual Machine (JVM).

4.2 Quantitative Results

To help with reading the results presented in this section, Table 1 provides a brief summary of the main terminology used in the description of the experiments.

We begin our evaluation by measuring the sustainable performance of the Hand Tracker in its vanilla non-Java, non-RAPID implementation, when executed in the high-end server and the low-end laptop respectively. The results of this analysis are the baseline of our evaluation and are displayed in Figure 4. The high-end hardware available in the server computer allows the C++ native application to achieve real-time processing at a rate bigger than 40 frames per second. The 30 fps limit is important since it matches the rate at which the RGBD camera acquires new Depth and RGB frames. On the other hand, the slower laptop can only achieve an average rate of 13 fps for the same C++ implementation, which is much smaller.

Then, we proceed by measuring the performance of the RAPID-enabled implementation of the hand tracker, when executed on the server and laptop host respectively, without utilizing code offloading. The results of these experiments are again portrayed in Figure 4. This analysis enables us to identify the overhead introduced by wrapping the native code of the Hand Tracker inside a Java container using JNI. The results obtained at this step reveal the impact of data serialization, synchronization, and JVM overheads. We observe that in this configuration, the application’s performance is reduced when executed on the high-end server host. When executing on the low-end laptop host, where the potential GPU speedup is lower and the overall execution speed slower, the overhead introduced by the Java Wrapper is much less evident. The

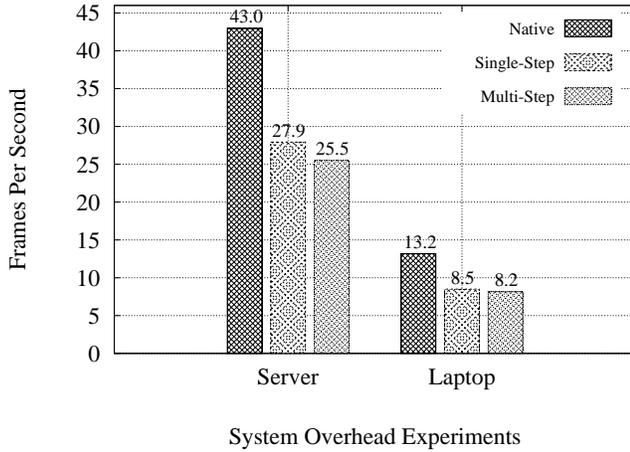


Figure 4: Performance graph of the native C++ implementation of the hand tracker in the two devices vs. the RAPID-enabled Java version without offloading. The server presents better results than the laptop, due to the outdated hardware of the laptop. The overhead added by the offloading framework is less pronounced in the laptop, due to the overall slower framerate achieved on the older hardware.

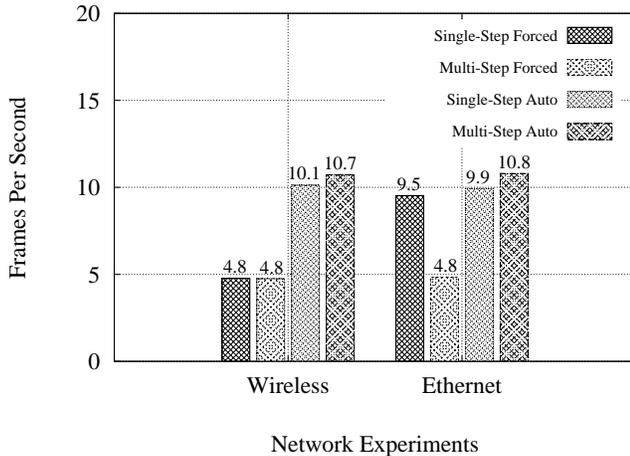


Figure 5: Performance graph of the offloaded hand tracking from the laptop to the server with Wi-Fi and Ethernet under different configurations. The *Forced* configuration presents better results when the hand tracking steps are implemented in a single Java method (*Single-Step*) and when Ethernet is used. The *Auto* configuration presents same results in all situations.

results also show that when wrapping each tracking step individually in Java methods (the *Multi-Step* case), the overhead is more visible compared to having all the steps in a single Java method (the *Single-Step* case).

Table 1: Terminology used in the experimental setup and related results.

Server	High-end machine with a GeForce GTX 1080M GPU and an Intel Core i7 processor.
Laptop	Thin-client laptop with an outdated GeForce 670M GPU and an early generation Intel Core i5.
Native	The C++ hand tracking implementation.
Single-Step	The hand tracking steps presented in Figure 2 are called inside a single Java method.
Multi-Step	The hand tracking steps presented in Figure 2 are called via separate Java methods.
Forced	RAPID is instructed to always offload an offloadable method. This is used to measure the performance in the case of a thin-client without GPU, which needs to always offload.
Auto	RAPID decides automatically whether to offload or not an offloadable method. This is used when the thin-client can run the method also locally, because it has a GPU.

Finally, we measure the performance characteristics of the distributed hand tracker, utilizing code offloading via the Gigabit Ethernet and the Wi-Fi connection. The purpose of this study is to evaluate the performance gain obtained by offloading the heavy parts of the application logic from the low-end laptop to the high-end server host.

When RAPID is instructed to always offload the offloadable methods (the *Forced* case), which would happen in the case of a device without a GPU, the *Single-Step* implementation under the Ethernet connection yields the best performance, with a framerate around 10 fps. This means that thanks to offloading, a machine without a GPU is possible to run the real-time 3D hand tracking with 1/3 of the desired framerate. Even though this might not be acceptable for a commercial use, it is still a good start towards building better systems that can improve these results.

If the thin-client is able to run the tasks also locally, such as in our case that the laptop has a GPU, RAPID can dynamically decide for each task to run it either locally or remotely. These results, which are noted as *Auto* in the graphs of Figure 5, show that RAPID is able to adapt in all situations and yield the best possible performance even if the connection is Wi-Fi rather than Ethernet. Also in this case, the achieved framerate is around 10 – 11 fps.

5 CONCLUSIONS AND FUTURE WORK

In this work, we presented the first results of the successful adaptation of a GPU-based real-time 3D hand tracker for executing it in thin-clients thanks to the computation offloading paradigm. We showed that, thanks to the maturity of the existing offloading frameworks, the effort of transforming the C++ native application to a distributed version was minimal.

Our experiments showed that choosing a Java-based offloading framework might not be ideal when it comes to real-time or near real-time applications. Indeed, our evaluations confirmed that the overhead of the Java layer is not negligible, and it considerably

reduced the performance of the application in terms of framerate. However, our experiments proved that Edge Computation offloading can help limited low-power devices execute applications that cannot be run otherwise because of lack of resources, such as the GPU in our hand tracking case.

As future work, we intend to work on two axes to improve the results presented here: *i)* investigate on the possibility of using or developing a real-time oriented offloading framework, so that applications will not be penalized by its overhead, and *ii)* investigate single-frame 3D hand pose estimators such as the one described in [33]. This could allow for a higher level of parallelization, since there would be no inter-frame dependencies. In such case, all newly acquired frames could be submitted directly to the computing resources without any stall, the network delay would not be accumulated, and offloading could provide a substantial improvement of the resulting computational performance.

ACKNOWLEDGMENTS

The work was partially supported by the European Commission Horizon 2020 programs RAPID (H2020-ICT-644312) and Co4Robots (H2020-ICT-2016-1-73186).

REFERENCES

- [1] S. Yuan, G. Garcia-Hernando, B. Stenger, G. Moon, J. Y. Chang, K. M. Lee, P. Molchanov, J. Kautz, S. Honari, L. Ge, J. Yuan, X. Chen, G. Wang, F. Yang, K. Akiyama, Y. Wu, Q. Wan, M. Madadi, S. Escalera, S. Li, D. Lee, I. Oikonomidis, A. Argyros, and T.-K. Kim, "Depth-Based 3D Hand Pose Estimation: From Current Achievements to Future Goals," in *IEEE CVPR*, 2018. [Online]. Available: <http://arxiv.org/abs/1712.03917>
- [2] L. Ge, H. Liang, J. Yuan, and D. Thalmann, "3D Convolutional Neural Networks for Efficient and Robust Hand Pose Estimation from Single Depth Images," in *Proceedings IEEE CVPR*, 2017, pp. 1991–2000.
- [3] C. Wan, T. Probst, L. Van Gool, and A. Yao, "Crossing Nets: Combining GANs and VAEs with a Shared Latent Space for Hand Pose Estimation," *CVPR*, 2017. [Online]. Available: <http://arxiv.org/abs/1702.03431>
- [4] M. Oberweger and V. Lepetit, "DeepPrior++: Improving Fast and Accurate 3D Hand Pose Estimation," in *ICCV workshop*, vol. 840, 2017. [Online]. Available: <http://arxiv.org/abs/1708.08325>
- [5] X. Sun, Y. Wei, S. Liang, X. Tang, and J. Sun, "Cascaded hand pose regression," in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 07–12–June, 2015, pp. 824–832.
- [6] D. Tang, H. Jin Chang, A. Tejani, and T.-K. Kim, "Latent regression forest: Structured estimation of 3d articulated hand posture," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 3786–3793.
- [7] A. Makris, N. Kyriazis, and A. A. Argyros, "Hierarchical Particle Filtering for 3D Hand Tracking," in *Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2015 *IEEE Conference on*, 2015, pp. 8–17. [Online]. Available: <http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=7301343>
- [8] 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.
- [9] P. Krejov, A. Gilbert, and R. Bowden, "Guided optimisation through classification and regression for hand pose estimation," *Computer Vision and Image Understanding*, vol. 155, pp. 124–138, 2017.
- [10] T. Sharp, Y. Wei, D. Freedman, P. Kohli, E. Krupka, A. Fitzgibbon, S. Izadi, C. Keskin, D. Robertson, J. Taylor, J. Shotton, D. Kim, C. Rhemann, I. Leichter, and A. Vinnikov, "Accurate, Robust, and Flexible Real-time Hand Tracking," in *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems - CHI '15*, 2015, pp. 3633–3642. [Online]. Available: <http://dl.acm.org/citation.cfm?doi=2702123.2702179>
- [11] D. Xu, Y. Li, X. Chen, J. Li, P. Hui, S. Chen, and J. Crowcroft, "A survey of opportunistic offloading," *IEEE Communications Surveys Tutorials*, vol. PP, no. 99, pp. 1–1, 2018.
- [12] M. Satyanarayanan, V. Bahl, R. Caceres, and N. Davies, "The Case for VM-based Cloudlets in Mobile Computing," *IEEE Pervasive Computing*, 2009. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5210087>
- [13] E. Cuervo, A. Balasubramanian, D.-k. Cho, A. Wolman, S. Saroiu, R. Chandra, and P. Bahl, "Maui: Making smartphones last longer with code offload," in *Proceedings of the 8th international conference on Mobile systems, applications, and services*. ACM, 2010, pp. 49–62. [Online]. Available: <http://portal.acm.org/citation.cfm?id=1814441>
- [14] B. Chun, S. Ihm, P. Maniatis, M. Naik, and A. Patti, "Clonecloud: Elastic Execution Between Mobile Device and Cloud," *Proceedings of the sixth conference on Computer systems*, pp. 301–314, 2011. [Online]. Available: http://dl.acm.org/ft_gateway.cfm?id=1966473&type=pdf
- [15] S. Kosta, A. Aucinas, P. Hui, R. Mortier, and X. Zhang, "ThinkAir: Dynamic resource allocation and parallel execution in the cloud for mobile code offloading," in *Proceedings - IEEE INFOCOM*, 2012, pp. 945–953.
- [16] H. Flores, X. Su, V. Kostakos, A. Y. Ding, P. Nurmi, S. Tarkoma, P. Hui, and Y. Li, "Large-scale offloading in the internet of things," in *2017 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*, March 2017, pp. 479–484.
- [17] D. Chatzopoulos, M. Ahmadi, S. Kosta, and P. Hui, "Openrp: A reputation middleware for opportunistic crowd computing," *IEEE Communications Magazine*, vol. 54, no. 7, pp. 115–121, 2016.
- [18] R. Montella, G. Giunta, and G. Laccetti, "Virtualizing high-end gpgpus on arm clusters for the next generation of high performance cloud computing," *Cluster Computing*, vol. 17, no. 1, pp. 139–152, Mar. 2014. [Online]. Available: <http://dx.doi.org/10.1007/s10586-013-0341-0>
- [19] R. Montella, C. Ferraro, S. Kosta, V. Pelliccia, and G. Giunta, *Enabling android-based devices to high-end GPGPUs*, 2016, vol. 10048 LNCS.
- [20] 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 Computation*, vol. 29, no. 24, 2017.
- [21] M. Clerc and J. Kennedy, "The particle swarm-explosion, stability, and convergence in a multidimensional complex space," *Evolutionary Computation, IEEE Transactions on*, vol. 6, no. 1, pp. 58–73, 2002.
- [22] S. Stefanou and A. A. Argyros, "Efficient scale and rotation invariant object detection based on hogs and evolutionary optimization techniques," in *Advances in Visual Computing (ISVC 2012)*. Springer, Jul. 2012, pp. 220–229.
- [23] P. Paderleris, X. Zabulis, and A. A. Argyros, "Head pose estimation on depth data based on particle swarm optimization," in *IEEE Computer Vision and Pattern Recognition Workshops (CVPRW 2012)*. IEEE, Jun. 2012, pp. 42–49.
- [24] I. Oikonomidis, N. Kyriazis, and A. A. Argyros, "Tracking the articulated motion of two strongly interacting hands," in *IEEE CVPR*. IEEE, Jun. 2012, pp. 1862–1869.
- [25] N. Kyriazis and A. A. Argyros, "Scalable 3d tracking of multiple interacting objects," in *IEEE Computer Vision and Pattern Recognition (CVPR 2014)*. Columbus, Ohio, USA: IEEE, June 2014, pp. 3430–3437.
- [26] P. Panteleris, N. Kyriazis, and A. A. Argyros, "3d tracking of human hands in interaction with unknown objects," in *British Machine Vision Conference (BMVC 2015)*. Swansea, UK: BMVA, September 2015, pp. 123–1.
- [27] A. Qammar, N. Kyriazis, and A. A. Argyros, "Boosting the performance of model-based 3d tracking by employing low level motion cues," in *British Machine Vision Conference (BMVC 2015)*. Swansea, UK: BMVA, September 2015, pp. 144–1.
- [28] T.-H. Pham, A. Kheddar, A. Qammar, and A. A. Argyros, "Towards force sensing from vision: Observing hand-object interactions to infer manipulation forces," in *IEEE Computer Vision and Pattern Recognition (CVPR 2015)*. Boston, USA: IEEE, June 2015, pp. 2810–2819.
- [29] J. Vijay, E. Trucco, and S. Ivekovic, "Markerless human articulated tracking using hierarchical particle swarm optimisation," *Image and Vision Computing*, vol. 28, no. 11, pp. 1530–1547, 2010.
- [30] D. Michel, C. Panagiotakis, and A. Argyros, "Tracking the articulated motion of the human body with two rgbd cameras," *Machine Vision Applications*, vol. 26, no. 1, pp. 41–54, 2015.
- [31] A. Qammar, D. Michel, and A. A. Argyros, "A hybrid method for 3d pose estimation of personalized human body models," in *IEEE Winter Conference on Applications of Computer Vision (WACV 2018)*. IEEE, March 2018. [Online]. Available: http://users.ics.forth.gr/argyros/res_personalizedHumanPose.html
- [32] RAPID, "Distributed Hand Tracker Github Repository," 2018. [Online]. Available: <https://github.com/RapidProjectH2020/HandTrackerRAPID/>
- [33] P. Panteleris, I. Oikonomidis, and A. A. Argyros, "Using a single rgb frame for real time 3d hand pose estimation in the wild," in *IEEE Winter Conference on Applications of Computer Vision (WACV 2018)*, also available at arxiv. IEEE, March 2018.