

HiveTracker: 3D Positioning for Ubiquitous Embedded Systems

Cédric Honnet

Sorbonne University, ISIR, CNRS
Paris, France
cedric@honnet.eu

Gonçalo Lopes

NeuroGEARS Ltd
London, United Kingdom
g.lopes@neurogears.org




Figure 1: Left: the HTC Vive tracker - Right: our HiveTracker miniaturization.

ABSTRACT

Recent advances in positional tracking systems for virtual and augmented reality have opened up the possibilities for ubiquitous motion capture technology in the consumer market. However, for many applications such as in performance art, athletics, neuroscience, and medicine, these systems remain too bulky, expensive, and limited to tracking a few objects at a time. In this work, we present a small wireless device that takes advantage of existing HTC Vive lighthouse tracking technology to provide affordable, scalable, and highly accurate positional tracking capabilities. This open hardware and open software project contains several elements, and the latest contributions described in this paper include: (1) a characterization of the optical distortions of the lighthouses, (2) a new cross-platform WebBLE interface, and (3) a real-time in-browser visualization. We also introduce

new possibilities with an adaptive calibration to estimate transformation matrices of lighthouses, and an FPGA approach to improve precision and adaptability. Finally, we show how the new developments reduce setup costs and increase the accessibility of our tracking technology.

CCS CONCEPTS

• **Human-centered computing** → **Interaction devices.**

KEYWORDS

Indoor Positioning, Open-source, Low Cost, Wireless Device, Virtual Reality, 3D Tracking

ACM Reference Format:

Cédric Honnet and Gonçalo Lopes. 2019. HiveTracker: 3D Positioning for Ubiquitous Embedded Systems. In *Adjunct Proceedings of the 2019 ACM International Joint Conference on Pervasive and Ubiquitous Computing and the 2019 International Symposium on Wearable Computers (UbiComp/ISWC '19 Adjunct)*, September 9–13, 2019, London, United Kingdom. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3341162.3349295>

1 INTRODUCTION

Robust real-time positional tracking systems have been the target of research for decades to support applications in natural user interfaces (NUI), virtual reality (VR) and augmented

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).

UbiComp/ISWC '19 Adjunct, September 9–13, 2019, London, United Kingdom

© 2019 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-6869-8/19/09.

<https://doi.org/10.1145/3341162.3349295>

reality (AR). All of these interface paradigms require an accurate sampling of the spatial environment of a user, and have to be responsive in real-time in order to sustain the illusion of control. Recent developments in motion sensing and tracking technology for VR have finally moved the field of motion capture technology away from expensive custom engineered spaces with multiple cameras. The need to deploy VR in the consumer market has pushed for scalable and cost-effective solutions, mostly through inside-out computer vision based simultaneous localization and mapping (SLAM) [4, 8] or laser-based lighthouse approaches [1].

Because each device is taking care of its own localization, such inside-out solutions are inherently scalable, allowing for a much larger number of devices to coexist in the same space. This has triggered many researchers to start rethinking the entire approach to full-body motion tracking, replacing the passive reflective beads with active inside-out devices [3]. Unfortunately, size and accuracy remain a limitation for such motion capture applications. Inside-out SLAM based computer vision devices rely on cameras and expensive processors, which both make existing controllers more bulky and increase power requirements.

An alternative solution to inside-out tracking is to use a beacon-based approach, such as the one used in the HTC Vive Lighthouse tracking system. In this system, a lighthouse base station periodically sweeps the space with a laser at a fixed frequency. Photodiodes placed in tracked objects can detect these sweeps and using either information embedded in modulated light at the beam, or synchronization pulses, can reconstruct the exact angle at which the laser hit the object. Positioning information can be accurately recovered from the photodiodes as long as the timing of each laser hit can be sensed fast enough.

2 CONTRIBUTIONS

In previous work, we introduced the first HiveTracker¹ prototype [7], which takes advantage of a rare dedicated real-time and parallel processing feature in the nRF52 microcontroller made by Nordic Semiconductor to sense the lighthouse light sweeps, without requiring a dedicated FPGA. This allowed us to build an embedded device with up to five photodiodes that is small, accurate, and cheap enough for full-body motion tracking at scale.

In this paper, we describe the further development of our miniaturized prototype (see Fig. 1), and the main hardware and software challenges we had to overcome to increase the accessibility and ease of use of the HiveTracker beyond a proof-of-concept.

Optical distortion characterization

To better understand the limitations of our system, we designed a simple test version to characterize it and improve it accordingly.

Test jig. The 1st step was to verify that we can measure a coherent 10 mm translation wherever we are in our working zone. We used a CNC² to sample positions along a linear path. We moved the CNC on 7 locations, and for each of them, we sampled 30 position data every 10mm. The HTC Vive positioning being accurate with the HTC recommended setup, we explored the limits of its capabilities. Instead of making the base stations face toward each other, we made them look in the same direction, as represented in figure 2.




Figure 2: Characterization jig: each numbered square is a position of the CNC.

Result. As we can see in the figure 3, the HTC Vive positioning system is not linear in this configuration. The observed distortion shows that measuring a distance of 300mm near the base stations is not the same as measuring the same distance further away. Nevertheless, now that we have characterized it becomes possible to calibrate a distortion map so we can improve our positioning measurements later.




Figure 3: Characterization result: each numbered line is a measurement of a 300mm distance.

¹<https://hivetracker.github.io>

²XY-plotter by Makeblock: <http://www.makeblock.com>

Firmware

The entire development stack of the project is open source, with the goal of making all components as accessible as possible. For instance, we worked hard to make our firmware³ Arduino compatible. Fortunately, the latest Arduino "Nano 33 BLE"⁴ uses the same microcontroller (nRF52) as our prototype, which will allow us to ensure long term support.

Bluetooth Low Energy communication

The nRF52 microcontroller provides a proprietary radio communication protocol for high-bandwidth communication, but also a BLE interface. BLE communication is attractive due to its low power requirements and growing acceptance as a standard in web development stacks. In our goal to make setting up and interfacing with HiveTracker devices more accessible, we decided to embed a lightweight BLE protocol in the device firmware, allowing direct pairing from any WebBLE enabled browser.

As mentioned in the previous section, all our tools are chosen with the goal of making the project accessible and reproducible. Up to now WebBLE was the most easily accessible route for reading BLE data streams, but alternative approaches have been recently published by the ITP-NYU community: p5.ble.js⁵. These should hopefully help improve the accessibility of this project further.

Real-time visualization in a web app

To enable rapid-prototyping, visualization and testing of the HiveTracker performance, we developed an in-browser JavaScript interface, visualization and simulation framework using the Three.js library⁶. This cross-platform approach allows visualizing the 3D positions in any device running Chromium or Chrome (see figure 4).




Figure 4: Web app showing the 2 base stations (cubes) and the tracker (sphere).

³Firmware Repository: <http://github.com/hivetracker/firmware>

⁴Arduino Nano 33 BLE: <https://store.arduino.cc/usa/nano-33-ble>

⁵p5.ble.js by ITP-NYU: <https://ltpnyu.github.io/p5ble-website/>

⁶Three.js library: <https://threejs.org/>

This interface can be tested, and modified from our repository⁷. This will enable mobile applications such as smartphone web apps for motion capture or augmented reality.

3 DISCUSSION

Adaptive calibration and 3D reconstruction

For each photodiode laser hit, we can reconstruct the positioning information of the sensor relative to the origin of the base station. However, it is often desirable to position the device relative to a common reference frame. In our previous work [7], we directly used the SteamVR calibration procedure to extract the location of each base station in absolute VR world coordinates. This was useful when integrating with existing VR applications running on the computer, for example to keep the coordinate frame of the device registered with a VR headset. However, the base station beacons are completely passive and only require a power supply, so it would also be an advantage to allow the calibration of a common reference frame independent of SteamVR, for situations where we are running a motion capture system in full-wireless mode, independently from a computer.

To calibrate our own reference, we use a HiveTracker device as the origin, with four photodiodes arranged in a known planar configuration (see figure 5).




Figure 5: Squared structure used for calibration.

For each pair of photodiode hits from the same base station, we recover two angles, from which we can reconstruct the line from the photodiode to the origin of each base station (Figure 6A). By projecting each line onto a fixed distance plane in front of the beacon, we recover 2D positions of points at the surface of the beacon plane (Figure 6B). Given we also know, by definition, the 3D coordinates of each of these points on the origin device, we can fully define a set of 2D/3D correspondences ($N = 4$). This is an instance of the well-known Perspective-N-Point problem for planar fiducial markers, for which there are a number of efficient solutions [2, 5]. After independently computing the pose of the origin

⁷Web app Repository: <http://github.com/hivetracker/hivetrackerjs>

device relative to both beacons, we can make use of bundle adjustment methods such as the Levenberg-Marquardt algorithm [6] to further improve the overall estimate of the beacon poses relative to the origin. In principle, we could extend this procedure to a grid of devices to further improve the estimated quality, as in [2].




Figure 6: Schematic of reconstruction procedure. (A): Relation of polar coordinates to cartesian coordinates, where θ and ϕ are the angles recovered from the device. (B): Projection of angles into fixed distance plane to generate 2D/3D correspondences.

Hardware possibilities using FPGAs

Until now, we focused on keeping the hardware⁸ simple, affordable, and scalable. However, a new version of the HTC LightHouse has been released which brings new features that both reduce the cost of base stations and enable for a larger tracking range using multiple beacons. This new version introduces the need for a very different sensing architecture using modulated light. Our nRF52 approach is not suitable for this, but fortunately, some FPGAs have been released which are both very small and well supported by the open source community⁹ such as the Lattice iCE40UL-1K. Its 1.4mm x 1.4mm package allows us to stay compact, and to support more photodiodes.

4 CONCLUSION

In this work, we presented a new scalable wireless 3D tracking device that we believe has the potential to open up new possibilities in the field of motion capture and behaviour studies, athletics performance, art, and medical diagnosis. The main motivation was to develop the entire stack of hardware and software on top of open and accessible standards,

⁸Online hardware design: <https://upverter.com/profile/hivetrapper/>

⁹Open Source FPGA Toolchain: <https://sybiflow.github.io/>

with the goal of maximizing dissemination of this state-of-the-art technology. The challenges have been numerous, and while much work remains to be done, here we showed that all of the remaining steps are dependent on current, immediately available technology. Our goal with this publication is to inspire others to join our development journey.

ACKNOWLEDGMENTS

The authors would like to thank Adam Kampff, Yvonne Jansen and Alexis Polti for their support offered during the development of this project. We would also like to thank Julien Mellet for his early work on reconstruction algorithms, Chinmay Pendharkar for his vital WebBLE contribution, Olivia Seow and Chris Davis for their design and media expertise, and NeuroGEARS Ltd. for the financial support as well as the provided help. This project was partially funded by the "Human Computer Interface Prize" of the Hack-A-Day conference. Finally, this work was also partially performed within the Labex SMART (ANR-11-LABX-65), supported by French state funds managed by the ANR under reference ANR-11-IDEX-0004-02.

REFERENCES

- [1] M. Borges, A. Symington, B. Coltin, T. Smith, and R. Ventura. 2018. HTC Vive: Analysis and Accuracy Improvement. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems*.
- [2] S Garrido-Jurado, R Munõz-Salinas, F J Madrid-Cuevas, and M J Marín-Jiménez. 2014. Automatic generation and detection of highly reliable fiducial markers under occlusion. *Pattern Recognition* 47, 6 (2014).
- [3] F. Jiang, X. Yang, and L. Feng. 2016. Real-time full-body motion reconstruction and recognition for off-the-shelf VR devices. In *Proceedings of the 15th ACM SIGGRAPH Conference on Virtual-Reality Continuum and Its Applications in Industry - VRCAI '16*. ACM Press.
- [4] G. Klein and D. Murray. [n.d.]. Parallel Tracking and Mapping for Small AR Workspaces. In *2007 6th IEEE and ACM International Symposium on Mixed and Augmented Reality*. IEEE.
- [5] V. Lepetit, F. Moreno-Noguer, and P. Fua. 2008. EPnP: An Accurate O(n) Solution to the PnP Problem. *International Journal of Computer Vision* 81, 2 (2008).
- [6] D. W. Marquardt. 1963. An Algorithm for Least-Squares Estimation of Nonlinear Parameters. *J. Soc. Indust. Appl. Math.* 11, 2 (1963).
- [7] D.R. Quiñones, G. Lopes, D. Kim, C. Honnet, D. Moratal, and A. Kampff. 2018. HIVE Tracker: A tiny, low-cost, and scalable device for sub-millimetric 3D positioning. In *Proceedings of The 9th Augmented Human International Conference (AH2018)*. ACM.
- [8] T. Whelan, R. F. Salas-Moreno, B. Glocker, A. J Davison, and S. Leutenegger. 2016. ElasticFusion: Real-time dense SLAM and light source estimation. *The International Journal of Robotics Research* 35, 14 (2016).