HIVE Tracker: a tiny, low-cost, and scalable device for sub-millimetric 3D positioning

Darío R. Quiñones
Centre for Biomaterials and Tissue Engineering, Universitat Politècnica de València, Spain
daquco@doctor.upv.es

Gonçalo Lopes
Kampff Lab, Sainsbury Wellcome Centre, University College London
London, UK
g.lopes@ucl.ac.uk

Danbee Kim
Kampff Lab, Sainsbury Wellcome Centre, University College London
London, UK
danbee.kim@ucl.ac.uk

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

David Moratal
Centre for Biomaterials and Tissue Engineering, Universitat Politècnica de València, Spain
dmoratal@eln.upv.es

Adam Kampff
Kampff Lab, Sainsbury Wellcome Centre, University College London
London, UK
adam.kampff@ucl.ac.uk

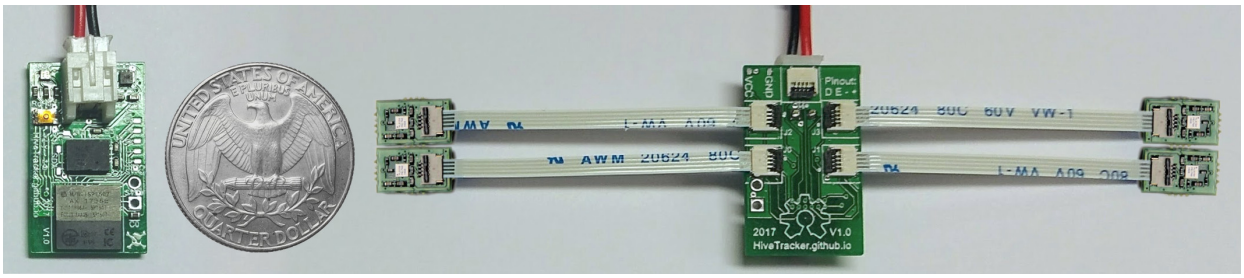

Figure 1: Hive Tracker prototype, with a US 25 cent coin for size comparison.

ABSTRACT

Positional tracking systems could hugely benefit a number of niches, including performance art, athletics, neuroscience, and medicine. Commercial solutions can precisely track a human inside a room with sub-millimetric precision. However, these systems can track only a few objects at a time; are too expensive to be easily accessible; and their controllers or trackers are too large and inaccurate for research or clinical use. We present a light and small wireless device that piggybacks on current commercial solutions to provide affordable, scalable, and highly accurate positional tracking. This device can be used to track small and precise human movements, to easily embed custom objects inside of a VR system, or to track freely moving subjects for research purposes.

AH2018, Feb. 7-9 2018, Seoul, Korea.

ACM ISBN 978-1-4503-5415-8/18/02…$15.00
https://doi.org/10.1145/3174910.3174935

CCS CONCEPTS

• **Human-centered computing** → **Virtual reality**; • **Computing methodologies** → **Motion capture**; • **Hardware** → **Wireless devices**; • **Software and its engineering** → **Open source model**; • **Information systems** → *Global positioning systems*;

KEYWORDS

Wireless-Sensor, Open Source, Virtual Reality, Motion Capture, Low cost, Indoor, Tracker, Neuroscience

ACM Reference Format:
Darío R. Quiñones, Gonçalo Lopes, Danbee Kim, Cédric Honnet, David Moratal, and Adam 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, New York, NY, USA, 8 pages. https://doi.org/10.1145/3174910.3174935

1 INTRODUCTION

Humans and other animals use bodies as our primary interface with the outer world, and as a powerful tool for expressing our inner worlds [1, 27, 28, 32]. Movement is therefore a phenomenon of interest to many, such as neuroscientists, surgeons, engineers, makers and artists. Though we have long appreciated movement on a macro scale, it has become increasingly clear that much could be gained from studying movement in greater detail and with higher precision.

1.1 Applications in performing arts and athletics

Two areas with long-standing interests in the precise study of human movement are the performing arts and athletics. In the 1920s, Rudolf Laban, a dance artist and theorist, developed with his colleagues a written notation system to precisely and accurately describe movement in terms of body parts, actions, floor plans, temporal patterns, and a three-dimensional use of space [17]. This has evolved into the modern-day Labanotation, or Kinetography Laban. However, many movement artists find Labanotation too complex and cumbersome for easy daily use. These days it is much more common to use videos to record, analyse, and teach the specific movements of performing artists and athletes. The increased prevalence of cheap video recording devices, especially phone cameras, has increased the use of this technique. But even high quality, state-of-the-art video recording technology struggles to capture the movements of aerial artists, acrobats, and other circus performers; the finer details of object manipulations performed by jugglers and athletes (e.g. finger placement and movements for optimal archery technique); and the tiny, fast movements of smaller body parts often found in hip hop and modern dance.

1.2 Applications in neuroscience and medicine

The fields of neuroscience and medicine are also interested in precisely recording and analysing movement, both for research and for clinical applications. A fundamental question in neuroscience is how nervous systems generate and execute movement goals. Historically, neuroscientists have prioritized the collection of cellular signals when answering research questions, so experiments are designed around the need to keep an electrode, or other data collecting device, stuck in the animal's head. This means that most neuroscience experiments study animals that are either held in place ("head-fixed") or tethered to a wire. However, a growing body of evidence supports the idea that cellular activity in the brain is significantly different when an animal is actively and freely moving in three dimensions through complex physical spaces, as opposed to when it is head-fixed or tethered to a wire. [10–14, 22]. Studies of the development and degeneration of nervous systems in humans also show that nervous systems are profoundly affected by the movements and physical contexts of their bodies [2, 18, 25]. Better systems for precise positional tracking in humans and animals would significantly impact the scientific questions that neuroscientists and clinical researchers could ask.

1.3 Current state-of-the-art in positional tracking

All of these areas would benefit hugely from having greater access to precise movement tracking. Motion capture systems, or mo-cap for short, are the current state-of-the-art for recording the movements of people and objects. However, current motion capture technologies require multiple specialized cameras, in addition to a whole slew of accessories, which unfortunately make these systems inaccessibly expensive and bulky. Industry standards for mo-cap, such as VICON and OptiTrack [3, 29], require a minimum investment of 10-15 thousand USD in order to assemble a viable system. Another option out on the market are inertial measurement units, or IMUs, which combine accelerometers, gyroscopes and magnetometers to track small and fast movements [8, 9, 30, 31]. Despite their sensitivity and speed, such IMUs do not measure absolute position, only relative motion. This means that to recover absolute position one must integrate the measured motion over time. Even small measurement errors will accumulate during this integration process, a phenomenon referred to as "drift". This makes IMUs inadequate for precise, continuous tracking of natural movement sequences. Another currently available alternative (at the time of writing) for motion tracking is Microsoft's Kinect, a consumer level 3D motion-sensing camera device for video game consoles. Using depth information and machine learning, the Kinect can infer in real-time the pose (position and orientation) of a human body located in front of the camera. However, a single Kinect cannot track motion in 360 degrees, as it was originally designed to track the movements of gamers facing a video game display. Some motion tracking systems combine Kinects and IMUs in an attempt to supplement one technology's weaknesses with the strengths of the other [24], but IMUs will always drift, and multiple Kinects will always be required to gain 360 degree tracking. While all these systems have found a niche, and are of great use to the military and entertainment industry, they are neither affordable to most artists, athletes, researchers, or clinicians, nor accurate enough for use in research or clinical settings [16].

1.4 Affordable, smaller, and more scalable

Simpler solutions are already coming out of artistic research, wearable applications and even implant experimentations [5, 19, 23]. Building upon this line of work, we present here an affordable, compact, and scalable positional tracking device called "Hive Tracker" (Figure 1), which can measure movement with sub-millimetric precision along six degrees of freedom; allows for untethered movement within a 5x5x5 m^3 space; connects easily and simply to virtual reality; and can scale up to as many devices as desired. This approach would allow the niches described above to take advantage of precise positional tracking technology, and opens the door to a plethora of new human augmentation applications.

2 MATERIALS AND METHODS

We present here the details of our multi-component system, composed of commercial products as well as custom devices and software. We first developed and benchmarked a small proof-of-concept of the system using an off-the-shelf microcontroller board described in Section 2.2.2. We present below some of this benchmarking data, which we used to constrain the subsequent design of the first Hive tracker prototype, described in Section 4.

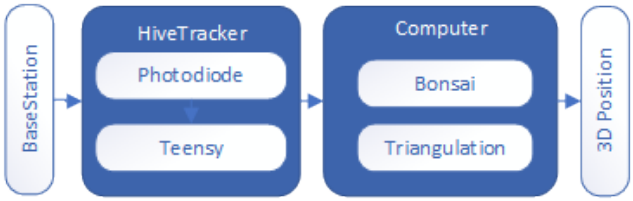

Figure 2: Overall system overview

Figure 2 shows an overview of the complete signal processing pipeline, which we describe in the following section.

2.1 Valve tracking system

The Hive Tracker is a data-collection device that piggybacks on a commercial virtual reality system developed by Valve (HTC VIVE). The commercial system consists of a headset, two hand controllers, and two light-emitting devices ("lighthouses" or "base stations"). Each lighthouse contains an LED matrix and two mirrors mounted on high precision rotors. In order to achieve sub-millimetric precision, this system needs to be set up in a space no larger than a 5x5x5 meter cube. We reverse-engineered the communication protocol between the lighthouses and the commercial devices in order to replace the commercial devices with custom devices optimized to fit our needs.

The signal from the lighthouses is composed of four components [Table 1], which enable the system to synchronize the two lighthouses and to track devices within the VR space (Figure 3). When the system initializes, the lighthouses are automatically assigned as lighthouse A and B. To synchronize the two lighthouses, lighthouse A first emits a flash of light with known pulse length. Soon after, lighthouse B emits a similar flash of light, as described in Table 1. The length of the flash determines which lighthouse will start the laser plane sweep, and whether that sweep will be horizontal or vertical. (Figure 4)

2.2 Signal processing

2.2.1 Photodiode circuit. In both our proof-of-concept and our first PCB prototype, we used the Chiclet, a sensor processing development board made by Triad Semiconductors. The first-generation of the Chiclet uses the TS3633 integrated circuit (IC). The IC converts a weak and noisy analog signal (obtained with the photodiode) to a digital signal which is simpler to use with a microcontroller. It provides both high-gain noise filtering and envelope detection of pulsed IR light that is incident on the photodiode.

Pulse start, μs	Pulse length	μs Source station	Meaning
0	65 - 135	A	Sync pulse
400	65 - 135	B	Sync pulse
1222 - 6777	~10	A or B	Laser plane sweep
8333	1556		End of cycle

Table 1: Activation Timings

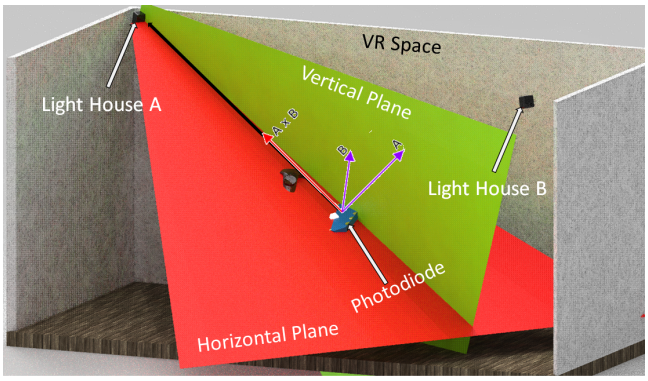

Figure 3: Representation of the computed lines in the 3D space. Picture represents how two intersected planes represent a 3D line pointing to each base station.

2.2.2 Acquisition Hardware: Teensy. The Hive Tracker proof-of-concept was developed on a Teensy 3.2 (PJRC.COM, LLC., Sherwood, Oregon, USA). This 35x17mm microcontroller uses an ARM[1] cortex M4 processor overclocked at 120MHz to reduce interrupt handling latency. The Teensy timestamps the digital signals coming from the TS3633 and sends them to a computer to be converted into angles (Section 2.2.3).

2.2.3 Acquisition Software: BONSAI. For data collection, integration, and online visualization, we used the Bonsai visual programming language [7]. Photodiode activation timestamps were collected and serialized into User Datagram Protocol (UDP) packets via the Open Sound Control (OSC) protocol [26]. These packets were streamed wirelessly into the host computer using WiFi. To reconstruct the position and orientation of each HiveTracker, we used the VR package of the Bonsai programming language to access the estimated 6 DOF (degrees of freedom) location of each lighthouse in parallel with the OSC messages.

2.2.4 Triangulation algorithm. For each light signal sequence (Table 1), each lighthouse will first flash, then scan a beam of light either horizontally or vertically [21]. Every photodiode gets hit by both the flash and the scans, but the light hits each photodiode at different times. Each lighthouse sweeps at 120Hz. The "incident plane" is the plane defined by the angle between a photodiode and a lighthouse (Figure 3). The cross product of the normals of the horizontal and vertical incident planes defines a vector ("incident line") between the tracking device and the lighthouse. The absolute position and orientation of each lighthouse is given by the

[1] Acorn RISC Machine - RISC: Reduced Instruction Set Computer

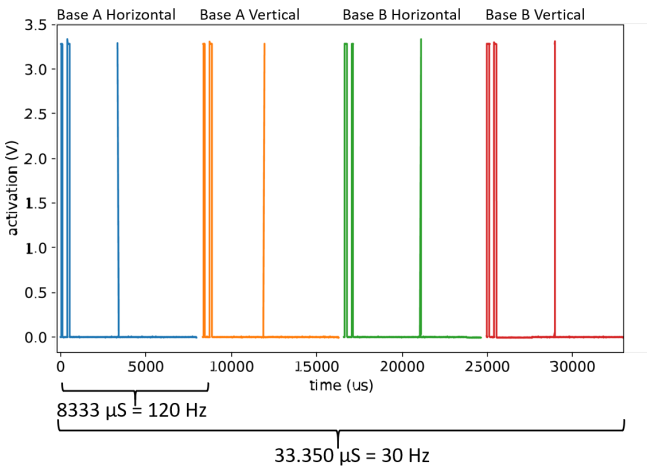

Figure 4: Light signal sequences (4 are shown): wide pulses are base station flashes and short pulses are laser plane scans.

commercial system, which allows us to project the incident lines from each lighthouse into the global coordinate system. The closest pair of points between these two incident lines defines the absolute location of a tracking device [4], which can be determined at 30 Hz (Table 1).

3 RESULTS AND DISCUSSION

3.1 Tracking inside an ideal room

We first compared the Hive Tracker proof-of-concept against the hand-held controllers of the commercial Valve tracking system in an ideal room. We taped a non-regular hexagonal shape on the floor of the testing room, then traced this shape by hand with both devices, recording the devices' positions using Bonsai. The acquired trajectories were overlapped to compare the accuracy of the commercial controller (32 photodiodes) against that of the first Hive Tracker proof-of-concept (1 photodiode). In this comparison, we used the commercial device as our baseline "ground truth". Since the tracing movements were parallel to the floor plane, we used only the sensors' X and Y axes for this benchmark (Figure 5).

To quantify the comparison, we fit a polygon shape to the tracking data from the commercial device. We then compared this trajectory to the average traces from the Hive Tracker by calculating the average distance of each point in the tracker trajectory to the fitted hexagon. The results of this comparison are shown in Figure 5. The Hive Tracker proof-of-concept, which uses only one photodiode, had an average error on the order of 10 mm more than the average error of a commercial tracker.

3.2 Tracking in a non ideal room

We then tested the Hive tracker proof-of-concept in the worst possible scenario: in a $1.2x1x1.5m^3$ space with a glass floor. In this situation the proof-of-concept was not able to achieve good accuracy, due to the short distances between the lighthouses and the reflection of the laser light on the glass floor (Figure 6).

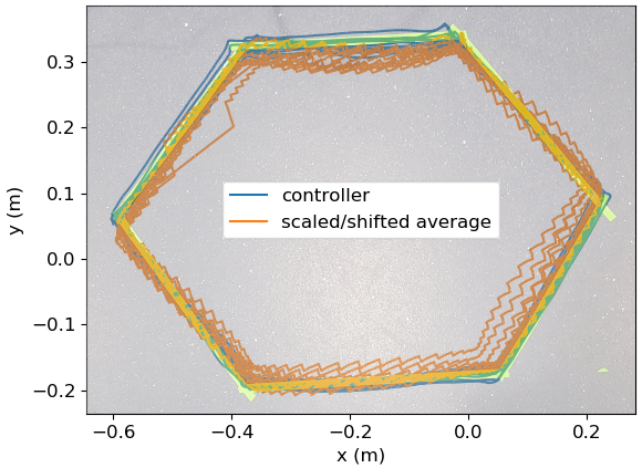

Figure 5: Accuracy comparison between a commercial controller and Hive Tracker proof-of-concept in an ideal room. Traces from both devices are superimposed over a photo of a polygon taped to the floor. Blue = commercial controller, Orange = Hive Tracker, Green = tape on the floor.

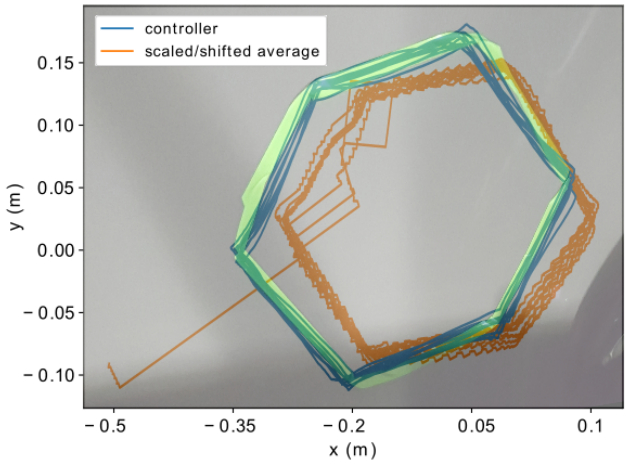

Figure 6: Accuracy comparison between commercial controller and Hive Tracker proof-of-concept in a non-ideal room. Traces from both devices are superimposed over a photo of a polygon taped to the floor. Blue = commercial controller, Orange = Hive Tracker, Green = tape on the floor.

3.3 Light reflections

Light reflections are a potential source of errors in this setup. As the laser plane (see Table 1) sweeps across the VR space, it is possible for light to bounce off a wall (or any shiny surface) and hit a photodiode sensor. Figure 7 shows false detections in the photodiode signal. We have addressed this issue in our first PCB prototype by adding extra photodiode sensors for more redundancy in the system. One can further shield the photodiodes from non-direct hits by embedding them in a shallow depression.

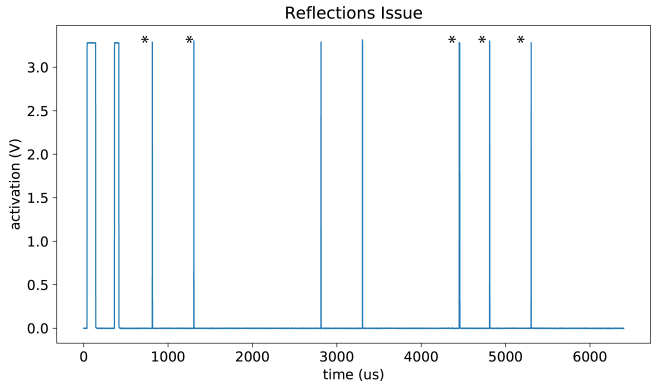

Figure 7: Reflections off the walls cause undesired pulses (asterisk) in the photodiode signal (compare to Figure 4).

3.4 Tracking refresh rate

Another limitation of the proof-of-concept was the refresh rate of positional measurements. In order to find the closest pair of points between the two incident lines from the lighthouses using only one photodiode, we needed to collect light data for at least four full cycles from each lighthouse. This means that the Hive Tracker proof-of-concept updates positional tracking measures at 30Hz (See Fig.4). This limitation can be addressed by using multiple photodiodes placed in a known geometric configuration relative to each other. In this way, the sequence of incident lines hitting each of the photodiodes in a single sweep can be used to constrain estimations of the absolute position and orientation of the Hive Tracker. This situation can be framed as a "Perspective-n-Point" problem, or the more general problem of estimating the position and orientation of a calibrated camera based on the projection of 3D points in the world onto that 2D camera image. Given this framing, we can consider each lighthouse as an ideal pinhole camera where the 2D-to-3D point correspondences are known exactly. Efficient algorithms to solve this problem for 3 or more points have been introduced by the computer vision community. [6]

We can further constrain the reconstruction by using inertial motion measurements to estimate motion over time. This makes it possible to reconstruct the position and orientation of an entire Hive Tracker device from single sweeps at 120Hz. These insights motivated the development of the first Hive Tracker PCB prototype described in the next section.

4 PCB PROTOTYPE AND NEXT STEPS

Our tests with the Hive Tracker proof-of-concept confirmed our initial suspicions, namely that we would need to use more photodiode sensors. However, this also increases the computational load on the microcontroller processing system, as multiple photodiode sensors may be hit simultaneously. To address this, our first PCB prototype includes a dedicated system for processing in parallel the photodiode signals. The hardware, firmware and software designs are open source and available online: http://HiveTracker.github.io

4.1 Hardware

Increasing the number of photosensors increases the computational load on the system's processing units. The most common way to deal with this is to use FPGA (Field Programmable Gate Array) processors, which enable true hardware parallel processing. However, we found a simpler approach to achieving the necessary parallelization while maintaining an extra-compact board, such that the device does not hamper free movement. We chose to use the nRF52 by Nordic Semiconductors (Oslo, Norway), a "System on Chip" (SoC) that replaces the functionality of the FPGA with a Programmable Peripheral Interconnect (PPI). The nRF52 also includes a BLE (Bluetooth Low Energy) and an ARM cortex M4 core.

Figure 8 shows the design of the first custom board, with 5 Chiclet connectors. On the top side of the board, shown on the right of the figure, the largest component (labeled "MCU" for Micro Controller Unit) is the ISP1507 by Insight SIP (Sophia Antipolis, France). This 8 x 8 mm^2 system-on-package (SoP) includes the nRF52, its necessary passives, a high accuracy crystal resonator to define the

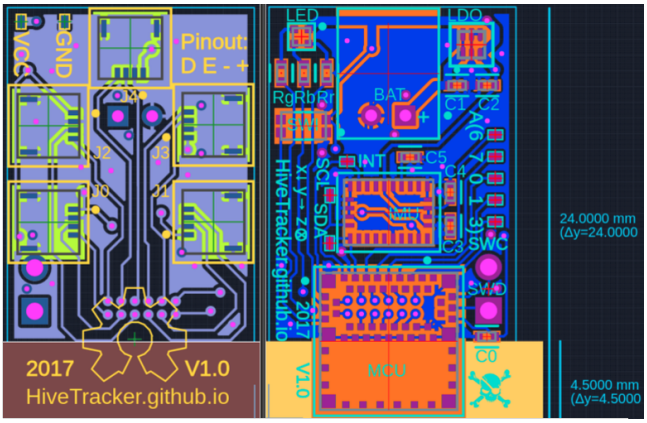

Figure 8: PCB design (left: bottom layer, with the photodiode connectors - right: top layer includes MCU, BLE, IMU, battery connector, button, LEDs, etc.)

radio communication speed, and a low power oscillator, which enables the microcontroller to save power while in deep sleep. The rectangular chip above the MCU is the Bosch BNO055 (Reutlingen, Germany) , an IMU SoC with a 3D accelerometer, a 3D gyroscope, a 3D magnetometer, and an ARM cortex M0 to perform sensor fusion. The other parts are the battery connector to connect LiPo batteries, a regulator, an RGB LED, a button, and 5 analog inputs that can behave as any GPIO (General Purpose Input/Output - to connect other MCU, sensors, etc). This first custom miniaturization prototype is far from cost-optimal, as the Chiclets are quite expensive. We have kept them in our design because their reliability is proven, and the cables connecting them to our PCB give us greater flexibility when placing the photodiodes. The next iteration of the Hive Tracker will not use Chiclets; instead, it will use the TS4231, a new IC by Triad Semi, and through-hole photodiodes that can accommodate custom orientations.

4.2 Firmware

The embedded software (firmware) configures the Triad Semi IC, processes the signal, merges it with the IMU data, and then sends it to a computer or a smartphone over BLE. This firmware fulfills the same function as the Teensy on the first proof-of-concept. The Teensy measures timing differences using interrupts, but this method can degrade the positioning accuracy. In the time it takes to handle an interruption, other light signals may have occurred. As mentioned earlier, an FPGA would solve this problem, but would make the Hive Tracker bulkier. While trying to keep the PCB small, we were able to validate that a rare feature of the MCU, the Programmable Peripheral Interconnect (PPI), could connect the edge detector and the capture register. This connection would normally need to happen via a CPU or FPGA, but using a PPI allows peripherals such as GPIOs and timers to interact autonomously with each other using tasks and events. The lack of interruptions makes it possible to simultaneously process up to 5 signals, which would improve robustness to potential occlusion. For a trackable object to detect IR signals in any 3D orientation, the geometric placement

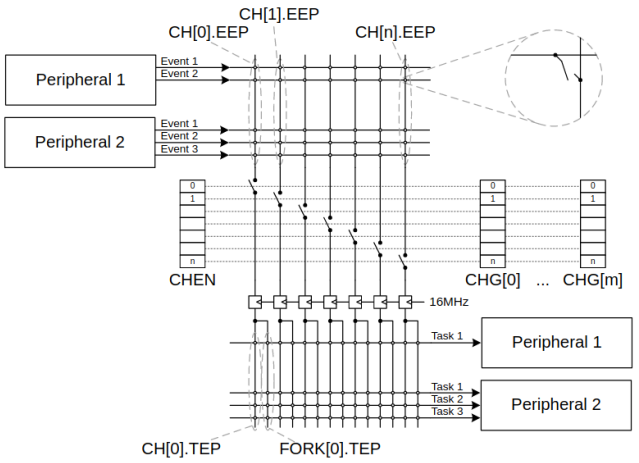

Figure 9: Programmable peripheral interconnect (PPI), ©Nordic Semiconductors [20]

of at least 4 photodiodes must be on the vertices of a tetrahedra. Figure 9 shows how these peripherals are connected.

This embedded software is developed using the SDK provided by Nordic Semiconductors, the MCU manufacturer, to allow optimal performances. It is also Arduino compatible, to allow the research community, makers, or artists to experiment openly and freely.

4.3 Size/accuracy trade-off

Our custom PCB prototype can fit into a bounding box volume about 100 times smaller than the one necessary to enclose a commercial tracking device, and our device weighs about 10 times less than the commercial tracker (Figure 10). These improvements in size and weight did cost some tracking precision (see section 3), but since our error calculations were made with the proof-of-concept, which used only one photodiode, we can treat those error measurements as a "worst-case scenario". The increased number of photodiodes in our PCB prototype can only improve the precision of the proof-of-concept.

These reductions in size and weight make the Hive Tracker much more convenient to use in a wider variety of human applications. Because it is not only small but also quite flat, the Hive Tracker can be integrated into clothing, gloves, or shoes worn by circus performers, theater artists, and dancers to capture their movements in real-time without hindering them. The Hive Tracker can also be easily integrated into objects manipulated by jugglers and athletes, in order to track and capture the movements of their props in addition to their bodies.

For medical and neuroscience research, which primarily use rodent subjects, the size and weight reductions are crucial. In neuroscience research, ethical review boards generally find it acceptable to use implants that do not exceed 10% of the body weight of the animal to be implanted. The average adult laboratory rat weighs between 250 to 500 grams [15], and the average adult rat head measures about 5 cm in length [15]. Given that the Hive Tracker measures 2.4 cm x 1.4 cm and weighs 8 grams, the Hive Tracker is already both small and light enough to be approved for use with rats, who are often implanted with devices that weigh 15 to 25 grams.

The accuracy of positional tracking that the Hive Tracker needs to achieve in order to be useful to medical and neuroscience research varies depending on the research question. Most research questions require a system that can precisely track the 3-dimensional movements of the whole body and the direction and orientation of the head. Current setups in neuroscience and medical research usually use video cameras to create a record of animal behavior, from which body trajectory and head movements are later extracted offline. This is both computationally and financially costly, and so many researchers simplify or ignore completely the behavioral validations required to thoroughly investigate their hypotheses. Even with the proof-of-concept's "worst-case scenario" precision, the Hive Tracker would already greatly increase researchers' ability to perform behavioral validations with a similar level of rigor as other controls currently used in neuroscience and medical research.

4.4 Cost

To produce the first 10 prototypes, the cost of the current version was about 75 USD, as opposed to 99 USD for the commercial tracker. This cost will be improved in our next version as the Chiclets won't be necessary anymore, so we can easily anticipate the cost of each Hive Tracker to be under 60 USD, especially if we produce them in larger quantities.

4.5 Autonomy

Various batteries can be used, but the one shown in figure 10 has a capacity of about 100mA. Given that our maximum power consumption is estimated to be about 40mA, this version of the Hive Tracker can run autonomously for at least 2 hours, depending on the usage. Embedded devices are in sleep mode most of the time, so the autonomy would greatly depend on the desired refresh rate. In addition, we do not need to process nor transmit data during periods of inactivity. For the applications mentioned in this paper, the inactivity rate might range from 10% to 90%, so the Hive Tracker could potentially run autonomously for 20 hours.

5 CONCLUSION

This paper showed an affordable and scalable custom positional tracking device that can integrate easily with a commercial consumer-level VR system. The Hive tracker is totally wireless and battery powered, which allows us to attach this small tracker to humans, animals, or even objects without interfering with natural movements and complex interactions. Furthermore, there is no limit to the number of trackers that can be used simultaneously. Based on the proof-of-concept presented here, we plan to use the Hive Tracker in a variety of applications, including neuroscience experiments of natural behavior; tracking and capturing object manipulations; 3D haptics in VR; and detailed and precise documentation of movements in artistic and clinical settings. These applications of the Hive Tracker can directly enhance our understanding of interactive movement and behaviour in humans and other animals. Thus, devices like the Hive Tracker are crucial to the kinds of research necessary for developing the next generation of human augmentation tools.

Figure 10: Size/weight comparison. Commercial tracker: 10 x 10 x 4.2 = 420cm^3; 85g. Hive Tracker: 2.5 x 1.5 x 1.1 = 4.13cm^3; 8g.

ACKNOWLEDGMENTS

The authors would like to thank Yvonne Jansen and Alexis Polti for their support and expertise in the making of the latest Hive Tracker prototype. The authors would also like to thank NeuroGEARS Ltd for the financial support and all the help that they provide. Darío R. Quiñones is supported by grant "Ayudas para la formación de personal investigador (FPI)" from Universitat Politècnica de València. Darío also acknowledges financial support from the Universitat Politècnica de València, through the "Ayudas para movilidad dentro del Programa para la Formación de Personal investigador (FPI) de la UPV". David Moratal acknowledges financial support by the Spanish Ministerio de Economía y Competitividad (MINECO) and FEDER funds under grant BFU2015-64380-C2-2-R. 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] Michael L. Anderson. 2003. Embodied Cognition: A field guide. *Artificial Intelligence* 149 (2003). https://doi.org/10.1016/S0004-3702(03)00054-7

[2] Nicholai A Bernstein. 1996. *Dexterity and its development.* https://doi.org/10.1080/00222895.1994.9941662

[3] Chien-Yen Chang, Belinda Lange, Mi Zhang, Sebastian Koenig, Phil Requejo, Noom Somboon, Alexander A Sawchuk, and Albert A Rizzo. 2012. Towards pervasive physical rehabilitation using Microsoft Kinect. In *Pervasive Computing Technologies for Healthcare (PervasiveHealth), 2012 6th International Conference on.* IEEE, 159–162.

[4] David H Eberly. 2006. *3D game engine design: a practical approach to real-time computer graphics.* CRC Press, New York and San Mateo, CA. 1040 pages.

[5] Rachel Freire, Cedric Honnet, and Paul Strohmeier. 2017. Second Skin: An Exploration of eTextile Stretch Circuits on the Body Video Figure. *Tei* 17 (2017), 653–658. https://doi.org/10.1145/3024969.3025054

[6] Vincent Lepetit, Francesc Moreno-Noguer, and Pascal Fua. 2008. EPnP: An Accurate O(n) Solution to the PnP Problem. *International Journal of Computer Vision* 81, 2 (19 Jul 2008), 155. https://doi.org/10.1007/s11263-008-0152-6

[7] Gonçalo Lopes, Niccolò Bonacchi, João Frazão, Joana P. Neto, Bassam V. Atallah, Sofia Soares, Luís Moreira, Sara Matias, Pavel M. Itskov, Patrícia A. Correia, Roberto E. Medina, Lorenza Calcaterra, Elena Dreosti, Joseph J. Paton, and Adam R. Kampff. 2015. Bonsai: an event-based framework for processing and controlling data streams. *Frontiers in Neuroinformatics* 9 (Apr 2015). https://doi.org/10.3389/fninf.2015.00007

[8] Sebastian Madgwick. 2010. An efficient orientation filter for inertial and inertial/magnetic sensor arrays. *Report x-io and University of Bristol (UK)* 25 (2010).

[9] S. O. H. Madgwick, A. J. L. Harrison, and R. Vaidyanathan. 2011. Estimation of IMU and MARG orientation using a gradient descent algorithm. In *2011 IEEE International Conference on Rehabilitation Robotics.* IEEE, 1–7. https://doi.org/10.1109/ICORR.2011.5975346

[10] B. L. McNaughton, S. J. Y. Mizumori, C. A. Barnes, B. J. Leonard, M. Marquis, and E. J. Green. 1994. Cortical Representation of Motion during Unrestrained Spatial Navigation in the Rat. *Cerebral Cortex* 4, 1 (1994), 27–39. https://doi.org/10.1093/cercor/4.1.27

[11] Edvard I. Moser, Emilio Kropff, and May-Britt Moser. 2008. Place Cells, Grid Cells, and the Brain's Spatial Representation System. *Annual Review of Neuroscience* 31, 1 (2008), 69–89. https://doi.org/10.1146/annurev.neuro.31.061307.090723

[12] Hisao Nishijo, Taketoshi Ono, Satoshi Eifuku, and Ryoi Tamura. 1997. The relationship between monkey hippocampus place-related neural activity and action in space. *Neuroscience Letters* 226, 1 (1997), 57–60. https://doi.org/10.1016/S0304-3940(97)00255-3

[13] J. O'Keefe. 1979. A review of the hippocampal place cells. *Progress in Neurobiology* 13, 4 (1979), 419–39. https://doi.org/10.1016/0301-0082(79)90005-4

[14] J. O'Keefe and J. Dostrovsky. 1971. The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Research* 34, 1 (1971), 171–175. https://doi.org/10.1016/0006-8993(71)90358-1

[15] George Paxinos and Charles Watson. 2007. *The Rat Brain in Stereotaxic Coordinates, 6th edition.* Academic Press.

[16] Alexandra Pfister, Alexandre M. West, Shaw Bronner, and Jack Adam Noah. 2014. Comparative abilities of Microsoft Kinect and Vicon 3D motion capture for gait analysis. *Journal of Medical Engineering & Technology* 38, 5 (Jul 2014), 274–280. https://doi.org/10.3109/03091902.2014.909540

[17] Valerie Monthland Preston-Dunlop and Susanne Lahusen. 1990. *Schrifttanz: a view of German dance in the Weimar Republic.* Princeton Book Company Pub.

[18] Frank Röhricht. 2009. Body oriented psychotherapy. The state of the art in empirical research and evidence-based practice: A clinical perspective. *Body, Movement and Dance in Psychotherapy* 4, 2 (Aug 2009), 135–156. https://doi.org/10.1080/17432970902857263

[19] Andreas Schlegel and Cedric Honnet. 2017. From Ordinary to Expressive Objects Using Tiny Wireless IMUs. (2017).

[20] Nordic Semiconductors. 2017. Programmable Peripheral Interface documentation. (31 Oct. 2017). http://infocenter.nordicsemi.com/topic/com.nordic.infocenter.nrf52810.ps/ppi.html

[21] Alexander Shtuchkin. 2017. DIY Position Tracking using HTC Vive's Lighthouse. (31 Oct. 2017). https://github.com/ashtuchkin/vive-diy-position-sensor

[22] J F Soechting and M Flanders. 1992. Moving in three-dimensional space: frames of reference, vectors, and coordinate systems. *Annual review of neuroscience* 15 (1992), 167–191. https://doi.org/10.1146/annurev.neuro.15.1.167

[23] Paul Strohmeier, Cedric Honnet, and Samppa von Cyborg. 2016. *Developing an Ecosystem for Interactive Electronic Implants.* Springer International Publishing, Cham, 518–525. https://doi.org/10.1007/978-3-319-42417-0_56

[24] Yushuang Tian, Xiaoli Meng, Dapeng Tao, Dongquan Liu, and Chen Feng. 2015. Upper limb motion tracking with the integration of IMU and Kinect. *Neurocomputing* 159 (Jul 2015), 207–218. https://doi.org/10.1016/j.neucom.2015.01.071

[25] Shoshanna Vaynman and Fernando Gomez-Pinilla. 2005. License to Run: Exercise Impacts Functional Plasticity in the Intact and Injured Central Nervous System by Using Neurotrophins. *Neurorehabilitation and Neural Repair* 19, 4 (2005), 283–295. https://doi.org/10.1177/1545968305280753

[26] David Wessel, Matthew Wright, and John Schott. 2002. Intimate Musical Control of Computers with a Variety of Controllers and Gesture Mapping Metaphors. *Proceedings of the 2002 conference on New interfaces for musical expression* (2002), 1–3. http://dl.acm.org/citation.cfm?id=1085213

[27] Margaret Wilson. 2002. Six views of embodied cognition. *Psychonomic bulletin & review* 9, 4 (Dec 2002), 625–36. http://www.ncbi.nlm.nih.gov/pubmed/12613670

[28] Margaret Wilson and Günther Knoblich. 2005. The Case for Motor Involvement in Perceiving Conspecifics. *Psychological Bulletin* 131, 3 (2005), 460–473. https://doi.org/10.1037/0033-2909.131.3.460

[29] Markus Windolf, Nils Götzen, and Michael Morlock. 2008. Systematic accuracy and precision analysis of video motion capturing systems - exemplified on the Vicon-460 system. *Journal of Biomechanics* 41, 12 (Aug 2008), 2776–2780. https://doi.org/10.1016/j.jbiomech.2008.06.024

[30] Kris Winer. 2017. 9 DoF Motion Sensor Bakeoff, GitHub. (31 Oct. 2017). https://github.com/kriswiner/MPU6050/wiki/9-DoF-Motion-Sensor-Bakeoff

[31] Kris Winer. 2017. Affordable 9 DoF Sensor Fusion. (31 Oct. 2017). https://github.com/kriswiner/MPU6050/wiki/Affordable-9-DoF-Sensor-Fusion

[32] Daniel M. Wolpert, Zoubin Ghahramani, and J. Randall Flanagan. 2001. Perspectives and problems in motor learning. *Trends in Cognitive Sciences* 5, 11 (2001), 487–494. https://doi.org/10.1016/S1364-6613(00)01773-3