

Demo: An Affordable and Easy-to-Setup Ground Truth System to Facilitate Localization Research

Lukas Furtner, Markus Schuss, Maximilian Schuh, and Carlo Alberto Boano

Institute of Technical Informatics, Graz University of Technology, Austria

E-mail: lukas.furtner@student.tugraz.at ; {markus.schuss, schuh, cboano}@tugraz.at

Abstract

We present L-Cube, a low-cost, open-source, and easy-to-deploy reference system based on the Lighthouse principle built to facilitate localization research. We showcase how L-Cube allows, after a simple setup procedure, to track the position of multiple tags.

CCS Concepts

• **Networks** → *Location based services*; • **General and reference** → *Experimentation*; • **Hardware** → *Sensor devices and platforms*.

Keywords


L-Cube, Lighthouse v2, HTC Vive, Crazyflie, OptiTrack, Base station, RTLS, Location awareness, IoT, Calibration procedure.

1 Motivation


The increasing demand for location-aware IoT applications has led to the rapid development of RF-based real-time location systems (RTLS). Ground truth systems are essential for testing such RTLS, as they provide accurate reference data (i.e., a benchmark of real-world positions allowing to evaluate the effects of signal interference, multipath propagation, and environmental factors on RF signals) against which the system’s performance can be evaluated. By comparing the RTLS output with the ground truth, one can identify errors, optimize algorithms, and ensure high accuracy.

Need for low-cost ground truth systems. Camera-based tracking systems such as OptiTrack Motion or QualiSys QTM have become the most prominent reference systems currently in use, offering a high update rate to test robotics applications. However, such systems are *very expensive*. For example, setting up an Optitrack on a 2x2 m area can cost ≈ 5000 € [5]. When testing localization systems on a larger scale (e.g., in the context of UWB-based systems deployed in multi-room indoor environments [4, 7]), camera-based tracking systems would be prohibitively expensive. There is hence a need for more affordable reference systems that, ideally, do not raise privacy concerns due to camera installations.

Need for scalable solutions. In addition to covering a physically-large scale, one may also need to track many devices simultaneously when testing a RTLS. To this end, a reference system also needs to be able to track an arbitrary number of devices. Camera-based systems commonly rely on special passive markers, which need to be attached to the object(s) to be tracked. Different marker configurations (i.e., physical placement of three or more markers in a unique location) are necessary to reliably distinguish between multiple objects. This is often very difficult when dealing with the tiny embedded devices used in RF-based localisation systems [3]. While active markers such as QualiSys’ Active Traqr can reliably track 100+ objects over long ranges, they are prohibitively expensive (≈ 470 € per active marker).



(a) Localization tile



(b) Tracking cube



(c) Base station

Figure 1: L-Cube’s key components: tiles (a), assembled tracking cubes (b), and off-the-shelf base stations (c).

Ease of deployment. Deploying a reference system on a large-scale can be a complex, tedious, and error-prone task. For example, the OptiTrack, which is relatively easy to calibrate, still needs the user to manually configure each individual camera’s focus and exposure settings. Moreover, special calibration tools (such as a calibration wand) are used to collect a large amount (thousands) of samples from all cameras to compute the correct room geometry and camera positions, or a calibration square to set the origin of a coordinate system. While such a calibration process is effective, the need for additional (and costly) calibration tools, and the time-intensive “wandering” phase are far from ideal.

Contributions. We present L-Cube, an open-source¹ ground-truth localization system based on the Lighthouse v2 base station [8]. L-Cube solely consists of affordable off-the-shelf hardware (HW) and allows a simple system setup and configuration, making it a flexible, easy-to-deploy reference system. With L-Cube, we aim to provide the research community with a useful tool to benchmark location-aware IoT systems.

2 Using Lighthouses as Base Stations

To design an affordable and easy-to-setup ground-truth system, we forego the use of cameras due to their drawbacks (see Sec. 1), and resort instead to Lighthouse v2 base stations [8].

Lighthouse tracking technology. In the mid-2010s, a new generation of consumer virtual reality (VR) headsets supporting inside-out tracking (i.e., the ability for devices to localize themselves using reference points in the environment) was introduced. As one such new *low-cost* system, the HTC Vive VR system promises sub-mm-level tracking on a room scale utilizing base stations (see Fig. 1c) also known as *Lighthouses* to serve as reference points – each costing ≈ 200 €. Designed for regular users wanting to play games with a VR headset and motion controllers, the HTC Vive’s *setup process* is remarkably easy: take a controller, point at the monitor, and trace the outside of the room. This process is remarkably similar to the “wandering” and origin point calibration of expensive camera-based tracking systems.

¹L-Cube’s schematics and code are available at: <http://iti.tugraz.at/L-Cube>.




Figure 2: Overview of L-Cube’s modular architecture.

The tracking mechanism of the HTC Vive was reverse-engineered and open sourced [6]. The base station periodically emits a reference pulse followed by a wide infrared beam sweeping the room alternating in the vertical and horizontal axes, which can be received by cheap photodiodes on a tracking device. By computing the difference in time between the pulse and the beam (while knowing how fast the beam sweeps), a device can compute its relative position to the base station. For 3D tracking, a device can triangulate its position based on the base station position and the information obtained by the beam reception on multiple diodes (*sweeping*) or by combining information from multiple base stations (*crossing*).

The first version of the HTC Vive used an early base station design (*Lighthouse v1*) and was limited to two base stations deployed in a single room of $\approx 3.5 \times 3.5$ m. To mitigate this, the *Lighthouse v2* [8] system – jointly designed by HTC and Valve, often referred to as SteamVR base station 2.0 – increases the range to 7 m, forgoes the synchronization pulse, and uses instead a modulated laser encoding the current position of the beam. As such, each beam can be individually attributed to up to 16 base stations while additional calibration information of the base station is also sent as part of the beams finally providing an affordable tracking system at *scale*.

The gap to be filled. Unfortunately, most open-source implementations of tracking devices are based on the Lighthouse v1 and are incompatible with the Lighthouse v2. Moreover, due to the high data rate of the modulated beam data (6 MHz), decoding and processing multiple sensors has so far only been achieved using costly FPGAs [8], while MCUs have only been shown to support a single photodiode [1] (i.e., they do not allow sweeping). An open-source tracking system based on the Lighthouse v2 is Bitcraze’s Crazyflie Lighthouse positioning deck [2]. Whilst intuitive to use and easy to setup, unfortunately, the localization process is deeply embedded into the Crazyflie’s flight controller, which bounds the user to use Crazyflie HW (≈ 300 € per tracking device).

3 L-Cube: Overview & Key Features

We hence design a new tracking device based on the Lighthouse v2. Instead of using costly special-purpose HW (e.g., a FPGA) to decode the beam information, we only rely on inexpensive off-the-shelf components. Furthermore, we leverage the Crazyflie’s configuration procedure, which simplifies and speeds up the deployment of the proposed tracking device.

Architecture. Fig. 2 depicts the architecture of our proposed reference system, L-Cube, which leverages Lighthouse v2 base stations (Fig. 1c). At the heart of L-Cube are *tracking cubes* (Fig. 1b), which consist of five independent *tiles* (Fig. 1a). This arrangement allows for the reception of the base stations’ beams from a wider range of angles. Each tile has four photodiodes, each one attached to one TS4231 light-to-digital converter which creates an envelope (light

is present) and data output (modulated data stream of the optical carrier). These two outputs are connected to an RP2040 microcontroller unit (MCU). Leveraging its programmable input/output (PIO) capabilities, the MCU decodes the beam’s information and estimates the position and orientation of the tile. Both localization methods (sweeping and crossing) are implemented and *dynamically* selected during runtime: this allows to use fewer base stations to cover an area (while dynamically switching to the higher accuracy crossing where available). While a standalone operation of each tile is feasible², on a tracking cube each connected tile forwards its estimate via I2C to a *cube controller* (an ESP32-S3 MCU) aggregating the positions of all tiles and presenting a combined position estimate to the user. Due to the high connectivity capabilities of the cube controller, data can be presented either via a serial connection or via more sophisticated connections, e.g., using Wi-Fi and MQTT.

Setup procedure. To set up L-Cube, a user sequentially places the tracking cube on a predefined pattern of four reference positions (e.g., the corners of a 1×1 m square). After a button press on each position, the L-Cube software running on a computer estimates the relative location and orientation of the visible base stations, ultimately generating a configuration of the setup. The latter can then be uploaded to all cubes, allowing for an easy deployment.

4 Showcasing L-Cube

We demonstrate L-Cube using two Lighthouse v2 base stations, a few tracking cubes, and a localization tile connected to a computer. First, we demonstrate the setup procedure described above. The configuration obtained is saved on the computer, which can then update the configuration of any other tile. Second, the configured tracking cubes can be freely moved within the testing area. A user interface (UI) shows the audience the position and orientation of the tracking cubes, as well as the position and orientation of the base stations. Furthermore, the UI indicates the localization method being used: with clear line-of-sight to both base stations, the cubes perform the more accurate crossing method, whereas they may fall back to the sweeping method when the beams from only a single base station can be received.

Acknowledgments

This work has been partially supported by the BMK, BMAW, and FFG, Contract No. 881844: “Pro²Future” (Products and Production Systems of the Future), within the ENHANCE-UWB project.

References

- [1] S. Alvarado-Marin et al. 2024. Lighthouse Localization of Miniature Wireless Robots. *IEEE Robotics and Automation Letters* 9, 10 (2024).
- [2] Bitcraze. 2024. Lighthouse Positioning System. [Online] <https://shorturl.at/9XKeo> – Last access: 2024-09-26.
- [3] B. Kempke et al. 2015. Polypoint: Guiding indoor quadrotors with ultra-wideband localization. In *Proc. of the HotWireless Workshop*.
- [4] D. Molteni et al. 2022. Cloves: A Large-Scale Ultra-Wideband Testbed. In *Proc. of the EWSN Conf.*
- [5] Optitrack. 2024. Build Your own Motion Capture System. [Online] <https://optitrack.com/systems/> – Last access: 2024-09-30.
- [6] D.R. Quiñones et al. 2018. HIVE Tracker: A Tiny, Low-cost, and Scalable Device for Sub-millimetric 3D Positioning. In *Proc. of the AH Conf.*
- [7] M. Schuh et al. 2022. First Steps in Benchmarking the Performance of Heterogeneous Ultra-Wideband Platforms. In *Proc. of the CPS-IoTBench Workshop*.
- [8] A. Taffanel et al. 2021. Lighthouse Positioning System: Dataset, Accuracy, and Precision for UAV Research. In *Proc. of the IEEE ICRA Conf.*

²One can use a serial interface to forward the position to a PC.