

OpenHPS: An Open Source Hybrid Positioning System

Maxim Van de Wynckel
Web & Information Systems Engineering Lab
Vrije Universiteit Brussel
1050 Brussels, Belgium
mvdewync@vub.be

Beat Signer
Web & Information Systems Engineering Lab
Vrije Universiteit Brussel
1050 Brussels, Belgium
bsigner@vub.be

ABSTRACT

Positioning systems and frameworks use various techniques to determine the position of an object. Some of the existing solutions combine different sensory data at the time of positioning in order to compute more accurate positions by reducing the error introduced by the used individual positioning techniques. We present OpenHPS, a generic hybrid positioning system implemented in TypeScript, that can not only reduce the error during tracking by fusing different sensory data based on different algorithms, but also make use of combined tracking techniques when calibrating or training the system. In addition to a detailed discussion of the architecture, features and implementation of the extensible open source OpenHPS framework, we illustrate the use of our solution in a demonstrator application fusing different positioning techniques. While OpenHPS offers a number of positioning techniques, future extensions might integrate new positioning methods or algorithms and support additional levels of abstraction including symbolic locations.

KEYWORDS

OpenHPS; hybrid positioning; open source; processing network; indoor positioning

1 INTRODUCTION

Determining the location of a person or asset is an important aspect of various human-computer interaction (HCI) solutions. Position tracking can be used to create autonomous vehicles, navigation systems and in context brokers such as CoBrA [7] to create implicit interactions based on the location and possibly other contextual information [35]. While we mainly rely on the Global Positioning System (GPS) to determine our location outdoors, many other positioning techniques exist that work both, indoors as well outdoors and sometimes even have been designed for completely different use cases as described later in Section 2.1. Each positioning technique has its advantages and disadvantages depending on the environment where it is being used. Hybrid positioning systems use the combination of different techniques and sensors to determine a more accurate position through sensor fusion [8].

One of the disadvantages with most commonly used indoor positioning methods such as Bluetooth Beacons or Wi-Fi access point positioning is the requirement of some form of calibration or training. With our proposed hybrid system called *OpenHPS*¹, the goal is to achieve hybrid positioning during both, the navigation or tracking (*online stage*) and between tracking, training and calibration of the used positioning techniques (*offline stage*). Unlike many of the existing frameworks, the goal of OpenHPS is to offer a layered

abstraction supporting a wide range of positioning techniques and fusion algorithms.

Using our knowledge on various positioning techniques and algorithms which are discussed in Section 2.1 and Section 2.2, we identified and analysed the actors of our system as highlighted in Section 4.1.1. This analysis revealed that actors play a different role depending on the positioning method being used. In order to allow the system to combine different existing positioning methods, we opted for a processing stream network where each node in a graph topology contributes to the sampled data.

Our design goal with the chosen processing network is to handle positioning data in real time, giving developers complete control over the data flow. The proposed OpenHPS framework is presented in Section 4 together with an illustrative demonstrator use case before discussing some future work in Section 5.

2 BACKGROUND

The first step in conceptualising our hybrid positioning framework was to investigate existing positioning methods in order to find their similarities as well as differences. OpenHPS has to be able to support a wide range of different positioning methods and implementation goals. When analysing existing approaches, we made a distinction between positioning *methods* and *algorithms*, similar to Wilson et al. [34]. Positioning methods represent the techniques and technologies that are available to determine a location while positioning algorithms include the algorithms that can be used to combine these methods.

2.1 Positioning Methods

In this section we list some of the more prominent positioning methods ranging from signal-based to visual solutions. Note that the selected positioning methods represent a limited set of the techniques and functional requirements we aim to support.

2.1.1 Global Position System (GPS). Starting with the most well-known technology, the Global Positioning System (GPS) is used for outdoor positioning [6]. It makes use of satellites in an orbit around the Earth to triangulate a location consisting of a longitude, latitude and elevation. There also exist variations on the original global position system, including Differential GPS (DGPS) [33] or the Assisted GPS (a-GPS) [13] that fuses GPS with dead reckoning described later in Section 2.1.5.

2.1.2 RF-based Positioning. RF-based positioning is a commonly used technique for indoor environments. Examples of technologies used for RF-based positioning include Wi-Fi, Bluetooth Beacons, RFID and LTE cell towers. These RF signals offer a *landmark* that

¹<https://www.openhps.org>

can be used as a reference when determining the position based on fingerprinting or other mathematical calculations.

2.1.3 Simultaneous Localisation and Mapping (SLAM). Simultaneous Localisation and Mapping, or SLAM for short, is a positioning method that makes use of 2D sensors to map its surroundings. One of the more common examples includes the use of LIDAR (Light Detection and Ranging) [28], capturing distance readings around a sensor. These readings can then be used to generate a 2D map of the environment.

2.1.4 Visual Positioning Techniques. Existing visual positioning techniques use image sensors to determine the position of the object the sensor is attached to (e.g. Visual SLAM [38]) or the position of objects in its field of view. In previously mentioned positioning methods, the *tracked* object obtains the sensor data. With Multi-Target Multi-Camera Tracking (MTMCT) [23] the tracked object is moving within the field of view of one or more image sensors.

2.1.5 Dead Reckoning. Dead reckoning calculates the current position based on the previous known position and the velocity or acceleration that is applied to that position [3, 20]. While the accuracy of dead reckoning is not ideal, it can be used to improved other positioning techniques such as GPS.

2.2 Positioning Algorithms

In order to calculate a position or to combine multiple positioning methods, different algorithms are used. The list of algorithms presented in this section offers a baseline for the types of positioning algorithms to be supported, but our OpenHPS framework is not limited to the presented set of algorithms.

2.2.1 Triangulation and Multilateration. Mathematical operations such as triangulation and multilateration can be used for several positioning methods and technologies discussed in Section 2.1. For techniques that provide a landmark such as RF-based positioning, the received signal strength (RSS) might provide a rough estimate of the distance. Other examples include mathematical positioning using a time difference with respect to the time of arrival (ToA) or the angle of arrival (AoA).

2.2.2 Fingerprinting. While with multilateration we only need information about the position of the used landmarks, fingerprinting requires a calibration for all possible positions in the tracking area [39]. A fingerprint of the sensor data at a given provided position is created during the offline stage. Later, these stored fingerprints are used during the online stage to reverse the sensor data into a position.

2.2.3 Noise Filtering. Sensor data should be filtered, which can be done through different noise filters. Similar to dead reckoning, a noise filter often requires knowledge of previous sensor readings and positions to predict the next result.

Noise filtering is one of the main requirements of our hybrid positioning system. The reason why we want to combine multiple technologies or algorithms is to reduce errors and noise filtering is the key component in realising this error reduction of positioning data. Note that individual positioning methods such as object

recognition or dead reckoning may want to perform different types of noise filtering algorithms tailored to their data.

2.2.4 Machine Learning. This type of algorithms includes a number of machine learning algorithms that can be of aid during calibration as well as positioning. These algorithms require training during the offline stage with the results being deployed during the online stage. This issue will be further discussed when discussing the requirements in Section 4.1, where we allow data to be used in the online and offline stages of the OpenHPS framework.

2.2.5 Computer Vision. With the visual positioning methods discussed in Section 2.1.4, the *tracked* actor is not always uniquely identified. Such visual positioning methods have to be able to detect and track objects between multiple frames, camera angles or positions. The algorithms used to track and detect persons or objects from a video stream are beyond the scope of our framework, but OpenHPS should be able to provide a generic interface supporting these types of algorithms.

2.3 Hybrid Positioning

Apart from supporting different positioning methods and algorithms, OpenHPS should be able to combine these methods. This requires a choice of algorithms to specify how the result of each method can be used in the combined output.

Sensor fusion can occur at a low or high level [14]. Raw sensor data such as IMU sensors or the relative signal strength from a transmitter can be fused in noise filtering algorithms. On a high level, calculated or provided positions (i.e. by third-party positioning systems) with a certain predicted accuracy can be combined using linear regression, heuristic weighted averages or any *decision fusion* algorithm.

3 RELATED WORK

Location-based Services (LBS) represent a generalised category of systems that provide the current location of a person or other objects [24]. A distinction between a push- and pull-based LBS is made [37]. A pull LBS provides a location when it is requested to do so, while a push LBS delivers information when a new location is determined by a provider.

The idea of combining multiple positioning methods in an LBS is not new. In this section we present some related work, ranging from existing hybrid positioning systems, frameworks, their used terminology and standards throughout location-based services.

3.1 Location APIs and Specifications

On the Web, the Geolocation API [31] offers a high-level interface for single or repeated position updates. The API provides an abstraction of the underlying technologies and algorithms used to determine the position. However, developers can request a high accuracy result or maximum cache age if hardware permits this. Resulting positions are geographical coordinates complying with the WGS84 standard [11].

JSR-179 and the improved JSR-293 [2] specifications are Java 2 Micro Edition (J2ME) modules that provide developers an API to obtain the location and orientation of a mobile device. Included in the API is a storage interface for landmarks (see Section 2.1.2). The

specification represents locations as timestamped coordinates with an orientation, accuracy, speed and information about the used positioning method [12]. When requesting a location, criteria such as the desired accuracy, power consumption and response timeout can be provided.

WebXR [22, 27] is a Web API that provides an interface for the tracking and use of VR or AR headsets. The API uses the *pose* terminology to indicate the position and orientation of the person wearing an XR headset in 3D space. While WebXR should not be considered as a location API, its specification uses terminology that is common in our framework. As an API that provides a tracking position, it adds a goal for our framework to support these third-party APIs.

3.2 Hybrid Positioning Systems

Various research concerning the fusion of sensor data to predict a more accurate position exists. SignalSLAM [29] represents an example of a hybrid system that uses signals of various positioning methods such as GPS, Wi-Fi and Bluetooth to map the surroundings. Chen et al. [8] have shown how a smartphone can combine sensor data of Wi-Fi access point positioning and Pedestrian Dead Reckoning (PDR). This combination of dead reckoning with another positioning method is a common combination used by many hybrid systems. LearnLoc [30] is a smartphone-centred positioning framework that uses fingerprinting algorithms (k-nearest neighbours algorithm) in combination with various sensor data available on a smartphone to provide power-efficient indoor positioning. The consideration of power efficiency is a common requirement in mobile positioning systems. Other than achieving the most accurate position, these location-based services use sensor fusion to prevent the continuous use of precise sensors such as cameras or GPS.

IndoorAtlas provides a Platform-as-a-Service (PaaS) with a well-established Software Development Kit (SDK) for combining Wi-Fi, GPS, Bluetooth beacons, dead reckoning and even geomagnetic positioning [18]. While the latter method has been found to be less ideal in steel reinforced buildings [30], it still offers a useful addition for creating a hybrid positioning system where geomagnetic positioning might be combined with other positioning methods.

In the research by Bekkelien and Deriaz [4] a framework called Global Positioning Module (GPM) had been presented for in- and outdoor positioning. GPM provides a uniform interface to different position *providers*. These providers are fused in a *kernel* that selects the position based on provided *criteria* (e.g. precision, accuracy or detection probability). Their approach offers a clear methodology on how this criteria can contribute to the selection or fusion of different technologies. However, the position providers and kernels are implemented on a high level of abstraction providing no room for developers to choose different algorithms or fusion techniques.

Ficco and Russo [15] presented a technology-independent hybrid positioning middleware called HyLocSys. Position *estimators*, representing different technologies, provide positions when a user performs a *pull* of their current position. Sensor fusion combines these estimated positions into a final response. With the middleware being an extension of the JSR-179 specification presented earlier, these pull requests accept criteria such as the preferred response time and expected accuracy. Other than many frameworks that

only provide geographical positions, HyLocSys provides geometric, symbolic as well as hybrid location models. Symbolic locations represent abstract places such as buildings, floors and rooms that are relatively positioned to each other. A hybrid location can convert this symbolic location to a geometric position. Note that the paper does not discuss positioning technologies such as dead reckoning or SLAM that require periodic updates in order to keep an up-to-date position.

Scholl et al. [36] propose a system that uses a LIDAR scanner to determine the fingerprinting position. This is somewhat similar to our goal of using different positioning methods to support the offline stage.

The Robot Operating System (ROS) [32] is a structured communication layer that can be used to create autonomous robots. It focuses on the integration of various robotics aspects such as positioning, computing and hardware interfacing. ROS provides the concept of peer-connected nodes that perform computational tasks. These nodes represent interchangeable software modules that help to build a pipeline from sensory data to an output action. For positions and orientations, ROS uses the *pose* concept which contains both the position and orientation of a user.

Our framework should adhere to specifications such as WGS84 when working with geographical positions. However, unlike many of the related work discussed in this section, we also want to support non-geometric positions. The hybrid location model presented by Ficco and Russo [15] offers a good type of location, but is still heavily focused on geometric positions.

Positioning methods and algorithms are often represented under the term *providers* that are optionally combined via high-level decision fusion. In our framework, we want to separate providers into generic algorithms and positioning methods that can easily be switched. This not only allows for more extensibility, but also some low-level sensor fusion.

The Geolocation API, JSR-179 and HyLocSys allow for the specification of accuracy or other criteria when requesting a position. However, unlike high-level APIs that hide the underlying technologies, OpenHPS is aimed towards developers with an understanding of the available hardware and positioning techniques that influence the criteria.

The landmark storage of JSR-179 is a very useful addition to a positioning system, as it is a common requirement for many positioning techniques. The persistence of landmarks between the online and offline stage is an important requirement that is extended to fingerprinting information and cached position storage in OpenHPS. This persistence should allow us to interface with existing systems such as the Geolocation API that support both push-based position updates as well as retrieving the current (cached) position.

4 OPENHPS FRAMEWORK

In this section, we present and discuss the system design of our proposed OpenHPS hybrid positioning framework. After listing some general requirements in Section 4.1, we outline the overall architecture in Section 4.2. Next, we provide some general information on our chosen implementation and demonstrate the use of OpenHPS in Section 4.3 and Section 4.4.

4.1 Requirements

Based on existing positioning methods and algorithms discussed in Section 2.1 and Section 2.2, the following framework requirements have been derived. We start by specifying the actors of our system and motivate the use of a processing network where each node of the graph topology might represent one of these actors.

4.1.1 System Actors. After investigating different existing positioning techniques, we defined four actors in OpenHPS:

- **Tracked actor:** This is an actor that can be tracked during the online positioning stage. A *tracked actor* can be an end user or an asset that might optionally contain sensors to further support the tracking. Our main goal is to determine the most accurate position of this type of actor.
- **Tracking actor:** This type of actor is responsible for tracking a tracked actor. Note that for some positioning methods, the same actor might act as a *tracking actor* as well as a tracked actor. However, for positioning methods such as the visual object tracking introduced in Section 2.1.4, the tracking actor is represented by the camera while a tracked actor is the object that is being detected.
- **Calibration actor:** Some positioning methods require a calibration before the positioning method can be used. Unlike the tracked actor, the purpose of a calibration actor is to train and calibrate how the tracking actor will be used in the online stage of the system.
- **Computing actor:** The computing actor is responsible for providing the final position output by our system. This actor combines the data generated by one or more tracking actors about a tracked actor and processes the data by, for example, using one of the positioning algorithms described in Section 2.2.

These four actors represent the four main components within OpenHPS. By distinguishing between the *tracked* and *tracking* actor, the system is able to support the tracking of persons or objects that do not actively participate in the positioning process.

4.1.2 Functional Requirements. In the following, we list the minimal functional requirements for our OpenHPS framework.

- **Online stage positioning:** In order to perform hybrid positioning or sensor fusion, multiple (processed) sources need to be combined by using different algorithms.
- **Offline stage positioning:** Processed results can be used to calibrate positioning methods of another (online) stage.
- **Third-party frameworks:** Our framework needs to support third-party high-level positioning systems. These external systems might provide their own calculated position of a tracked actor that needs to be fused with the position determined by our framework. In addition, the identification of this tracked actor might differ between frameworks.
- **Environment mapping:** With the requirement to support positioning methods such as SLAM and VSLAM, the system does not only offer the possibility to output an absolute position, but might also create an environment map. Our solution should be capable of handling, storing and using this map to its advantage.

- **Decentralisation:** Our positioning framework should be able to combine the four different actors introduced in the previous section based on remote hardware. This requires the framework to work decentralised without requiring any centralised sensor fusion, which can be achieved by allowing multiple computing actors to work independently. However, developers should still be given the option to centralise certain parts of the system if needed.
- **Monotonicity:** Partial information from a source should result in a partial output. In the context of a positioning system, this means that a computing actor does not need the sensor data of all positioning methods to determine a position. This requirement also helps in the decentralisation and parallelisation of the framework.

4.1.3 Non-functional Requirements. The following non-functional requirements contributed to the final decision about the software language used for OpenHPS.

- **Availability:** Our solution has to be available on various platforms ranging from servers to embedded systems; also supporting the decentralisation functional requirement.
- **Performance and latency:** Throughput is an important criteria when processing streaming data. Input data such as video and audio streams needs to be processed in real time. The latency also indicates how long it takes for data to be used in computations. As our goal is to achieve an accurate current position, outdated sensor data is not relevant.
- **Modularity:** The framework should be modular with both, a low-level API and modules that can be added and removed based on the available sensors and concrete use cases. Developers should remain in control of the types of algorithms and the flow of data from producer to consumer.

4.2 Framework Architecture

In order to support the presented functional and non-functional requirements, we decided to build on a *stream-based positioning system* that takes various types of *input data* and processes this data to get the desired output. Data that is transmitted between nodes is encapsulated in so-called *data frames* that can contain sensory data as well as one or more *data objects* the sensor data applies to, and are described in detail in Section 4.2.2.

For the design of our process network, a number of existing stream- and layer-based frameworks such as Akka Streams [9] or TensorFlow [1] have been investigated. These frameworks solve similar issues and are further detailed in Section 4.2.1. Due to the fact that each node needs to be configured individually, the decision was made to investigate flow-based frameworks where each component of the stream network is added individually.

Unlike low-level data stream frameworks, OpenHPS focuses on data that is helpful for positioning. We offer a higher-level API for creating the network and data that is handled by the system. Concepts such as edges or ports that are often found in stream-based programming languages are abstracted and not directly accessible by developers. However, unlike other hybrid frameworks [4, 18], the stream processing is extensible enough to give developers the opportunity to modify the positioning methods along with the used algorithms.

We start by discussing our process network design that uses a graph topology similar to other stream frameworks. Next, we present the data frames, objects and positional data that are being handled by the network.

4.2.1 Process Network Design. The OpenHPS framework uses a process network to handle data. The data is processed and dynamically manipulated by multiple connected nodes in a predefined graph topology. In the following, we list our three main design goals for this network:

- (1) **Consistent data types:** Data that is being processed in the network should have a reliable type and content. We process `DataObjects` encapsulated in `DataFrames`, which provides a defined scope how generic parts of our network should handle information.
- (2) **Processing goal:** Processing has the goal of providing an absolute position for our tracked actors. With this goal we have a clear understanding how every computing actor contributes to the output.
- (3) **Producer priority:** The producer or tracking actor has the highest priority. Slow consumers or computing actors must not result in outdated sensor information. Rather, developers should be given the opportunity to control what happens with the overflow of information that cannot be processed timely.

Starting from the goal of producing up-to-date positioning information, we opted for a push-pull-based stream. Data can be dropped if its not relevant for determining a more up-to-date position. The monotonicity of our framework ensures that positions can be determined based on partial data.

Each node can be designed to accept both push and pull requests. Similar to reactive streams [16], push and pull actions are *promise*-based and can be executed asynchronously. If a node that receives a pull request cannot respond with a data frame itself, it will forward the pull request to its incoming node(s).

Different to a traditional pull that returns a *response*, we use the *push* terminology to indicate a response for a given pull. This behaviour and terminology is similar to Akka Streams [9], but unlike reactive streams where data can only be provided when there is a demand, there is no back pressure built into the stream itself. Using the push terminology for a pull response removes the ambiguity of a response arriving after an already existing push in the pipeline. It also enforces the design goal of producers having the highest priority, even if a producer only generates information when requested.

A regular node has a unique identifier and push/pull functionality for data frames. Each node can have $0 \dots n$ inlets or outlets. Our system consist of the following three subtypes of the regular node:

- **Source node:** A source node provides a specific data type. This can either be a push or pull node that pushes data frames when they are available (e.g. a camera recording at a fixed frame rate) or creates a new data frame when the downstream node asks for it via a pull request (e.g. triggering a Bluetooth scan). The source node merges *data objects* in the data frame with those that were previously stored via *data services*. This merging behaviour prevents the need for

feedback loops to gain knowledge on previously calculated positioning data.

- **Processing node:** A processing node is a higher-level interface for a regular node. It provides an abstraction on the push and pull functionality to simplify the creation of a processing function of either data frames or individual data objects.
- **Sink node:** An output node or sink node accepts a specific data type as output frame. Unlike processing nodes, this type of node will not push data to other nodes. Upon receiving a data frame, the data objects will be stored using a compatible data service. Once saved, an event is sent upstream to indicate that the processing of this frame and its contained objects is completed.

Extensions of these nodes, allowing for specific data flow shapes and common position processing nodes, are provided in our core component. Figure 1 shows an example of a *positioning model* that has a source node, a sink node and four processing nodes connected in a graph structure. This positioning model describes a configured computational model aimed for processing sensor and positioning data [25]. Similar to existing streaming or pipelining frameworks, the graph can contain data flow shapes that manipulate the flow of data frames. Examples of such shapes include, but are not limited to balance nodes, data frame chunking, debouncing and merging of data objects and their processed positions.

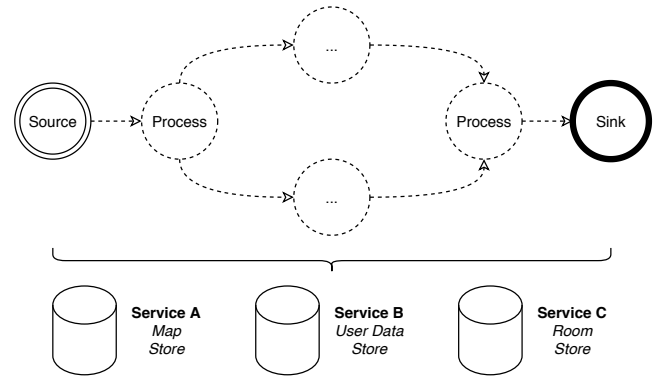


Figure 1: Example OpenHPS positioning model

All nodes in a positioning model have access to a set of services that allow the storage of objects. In the given example, three services are added for the storage of map, user data and room information. In our implementation, sink nodes always store data objects contained in received data frames. However, every node has the ability to fetch or insert new data into available services. This persistence allows for the storage of landmark objects, similar to the JSR-179 specification [12]. At the same time, these services can be used as an interface to fetch the latest position without requiring a specific implementation in the sink.

The positioning model can be created by using a builder pattern as illustrated in Listing 1. This builder creates the immutable properties of the model, including data services and the flow of data from source to sink. Models can have multiple flow shapes, each with one or more sources, processing nodes and sinks.

```

/* @openhps/core | version 0.2.0 */
1 ModelBuilder.create()
2   .addService(/* ... */)
3   .addShape(GraphBuilder.create()
4     .from(/* ... */)
5     .via(/* ... */)
6     .to(/* ... */))
7   .build().then((model: Model) => { /* ... */ });

```

Listing 1: Creation of a positioning model

In Figure 2, data is being pushed by an *active source* node. Processing nodes will process the data and push the modified frame to their output nodes. Push and pull actions are promise-based and resolved whenever the node finishes processing the frame. This allows for non-blocking asynchronous requests. The resolved push promise (indicated in green) gives an indication that the processing of the push is finished. However, it does not provide knowledge on whether or not the frame is processed by the complete network. To indicate this, sink nodes that receive a frame will emit a *completed* event that includes the data frame identifier and list of persisted object identifiers.

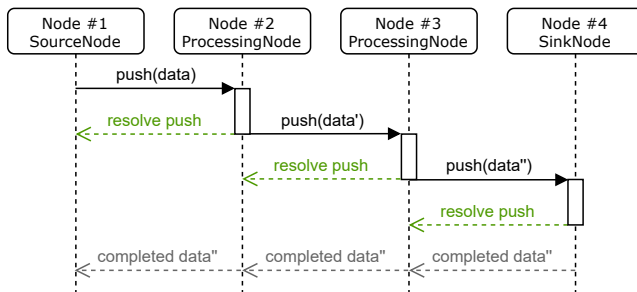


Figure 2: Data being pushed by a source

With the swim lane shown in Figure 3, the data is not automatically pushed by the source node. A downstream node such as a sink will send a `pull()` request to its input nodes. If these nodes cannot provide a frame of their own, the `pull()` request is forwarded to their respective input nodes. If the source has data available, a response to this pull is provided asynchronously. As mentioned in the beginning of this section, a `pull()` response will use the same invocation as a `push()`. In that case, the pull promise is resolved right after the source sends this push as indicated by the blue resolve chain in Figure 3.

As promises are resolved after the data frame is processed by a node, upstream nodes in the process chain cannot determine whether data has been processed successfully. Figure 5 shows a `push()` request that throws an error at the sink node (e.g. failure to store). An error event is triggered on previous node(s). By default, these nodes will chain the error to upstream nodes. However, each node can act upon this error in its individual implementation.

Nodes are implemented by developers on a high level of abstraction compared to other stream processing frameworks. Developers do not have the ability to push or pull from specific incoming or outgoing edges. Listing 2 shows two custom source nodes. The

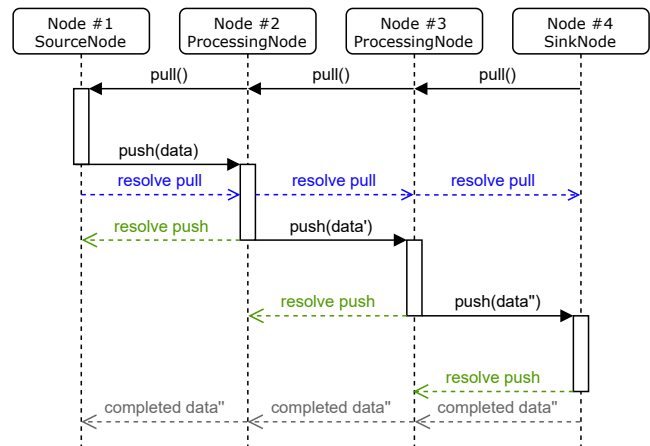


Figure 3: Data being pushed by source after receiving a pull

pull-based source node on lines 1 to 7 implements the `onPull()` function that is called whenever the source receives a `pull()` request. This function expects a promise of a data frame. Internally, the extended source node class will push this data frame as shown in Figure 3. With the push-based source (lines 9 to 22), the `onPull()` is unused. Instead, a timer is created that pushes a new data frame every 1000 milliseconds. A similar abstraction exists for sink nodes with the `onPush()` function.

```

/* @openhps/core | version 0.2.0 */
1 export class PullBasedSource extends SourceNode<DataFrame> {
2   public onPull(): Promise<DataFrame> {
3     return new Promise((resolve) => {
4       resolve(new DataFrame(this.source));
5     });
6   }
7 }
8
9 export class PushBasedSource extends SourceNode<DataFrame> {
10  constructor(source: DataObject) {
11    super(source);
12    this.on('build', () => {
13      setInterval(this._generate.bind(this), 1000);
14    });
15  }
16  private _generate(): void {
17    this.push(new DataFrame(this.source));
18  }
19  public onPull(): Promise<DataFrame> {
20    return Promise.resolve(undefined);
21  }
22 }

```

Listing 2: Push- and pull-based SourceNode classes

Similar to sources and sinks, processing nodes are abstracted. Any `pull()` requests to these nodes are automatically forwarded to the incoming nodes, as these process nodes do not generate new data frames. Developers are expected to implement a `process()` function manipulating a frame or individual objects within a frame.

4.2.2 Data Frame. Data that is pushed through the positioning model is represented within data frames, generated by a source

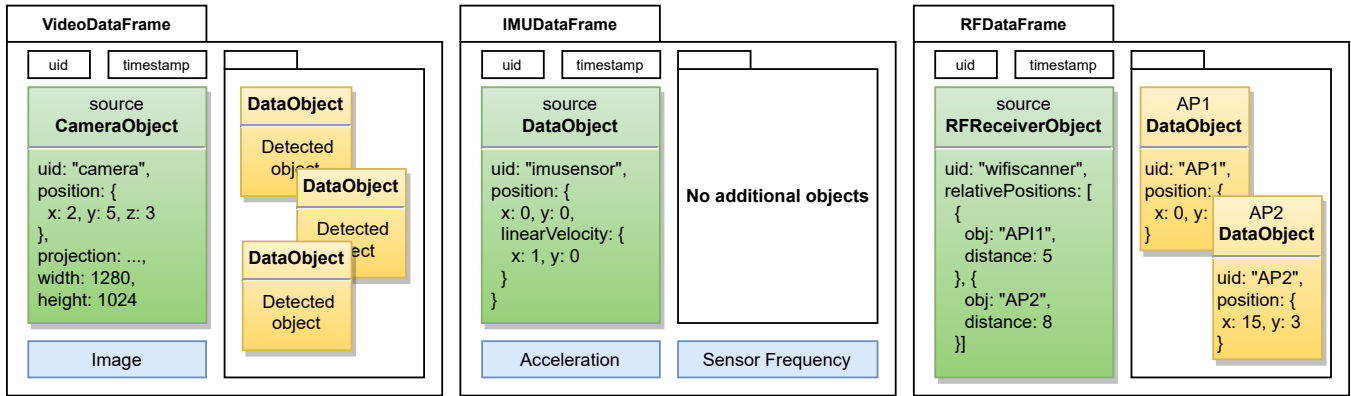


Figure 4: Data frame content examples

node. This ensures that the origin of data can be determined via some collection of metadata. The data contained in these frames includes (but is not limited to) the following attributes:

- **Unique identifier:** Each frame generated by a source is uniquely identified. This ensures that frames which are being processed by multiple processing nodes in parallel can be merged at a later stage in the stream.
- **Timestamp:** Required for determining when the data was created or obtained. When working with multiple sources that capture data of the same tracked actor, the timestamps will be used to merge the data frames. A timestamp is kept for the creation of each data frame by the source. This timestamp can also be used for time-based calculations such as applying velocity to a position. Using this timestamp instead of the system time results in a more deterministic output.
- **Source data object:** This is the data object that obtained the sensory data (e.g. the camera object or RF receiver). It is not always the object that is being tracked, but it can be required in order to determine the position of other objects (see actors in Section 4.1.1). Similar to the timestamp and identifier, the source data object can be used to specify certain criteria on how data frames or positions should be merged.
- **Data objects:** Data objects include everything that is of relevance to the positioning (e.g. the tracking and tracked actor). This also includes reference spaces needed for the positioning as pointed out later in Section 4.2.5. By grouping the data objects in the same data frame, nodes do not have to access any services to get this relevant information.

In order to demonstrate the content of data frames, Figure 4 depicts three situations where data is contained in frames. The first example shows a data frame created by a camera source. This camera object has a certain position and projection matrix. Linked to the data frame is a single image (i.e. video frame) captured by this source. During the processing of the image, objects can be detected and added to this frame before being pushed further downstream. In the second example we show data obtained by an accelerometer. The source object has a velocity and position, the frame itself contains the current acceleration and sensor frequency. This information can be used by a processing node to add the acceleration to the

existing velocity. In our third and final example, we show a data frame created by a Wi-Fi scanner. The scanner (source) has two relative distances to access points (AP). The information, mainly the position of these access points is included in the frame.

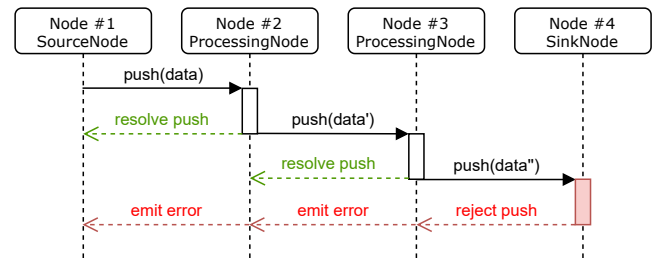


Figure 5: Error handling in push() request

4.2.3 *Position.* Similar to existing work [17, 26], OpenHPS distinguishes between *relative* and *absolute* positions. Absolute positions represent a fixed position in a specified *space* while relative positions indicate the position relative to another object. Absolute positions contain the following information:

- **Timestamp:** The time when the absolute position has been recorded or modified. The timestamp can be set by the sensor or by a processing node that calculated the position.
- **Accuracy:** General position accuracy with the same unit as the position itself. In the context of a hybrid positioning system, the accuracy can be used as a weight when merging with other calculated positions.
- **Orientation:** Stationary orientation of the data object at the recorded position. This orientation is relative to the X -axis and is represented in quaternions. However, it is possible to convert the quaternion representation to (and from) Euler or axis angles.
- **Linear velocity:** Linear velocity at the recorded position, relative to the orientation of the object (see Figure 6) using the axis X_{Obj} and Y_{Obj} of the point P with orientation ϕ .
- **Angular velocity:** Similar to linear velocity, the angular velocity is relative to the orientation of the object.

- **Position vector:** Each position can be converted to a three-dimensional vector, which enables the use of 2D positions in 3D reference spaces.
- **Unit:** Length unit of the position. This unit applies to the position vector and its accuracy.

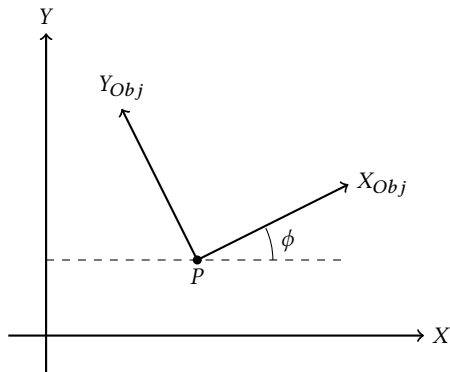


Figure 6: Position representation

Relative positions have the following attributes:

- **Timestamp:** Similar to the absolute position, this is the time when the relative position has been recorded.
- **Accuracy:** General position accuracy in a specified unit.
- **Reference object:** The referenced data object to which the position is relative to.
- **Reference value:** The value that determines the relative position to the reference object. This can be a distance, angle or velocity.

By default, OpenHPS and its core positioning algorithms support 2D, 3D and geographical coordinates. Developers can further extend the coordinate space with higher-level absolute and relative positions. Positions can be stored with a specified unit (i.e. length unit for absolute positions) in order to offer developers flexibility in the stored precision. Linear and angular velocity values are converted to a fixed unit (m/s for linear and rad/s for angular velocity). However, this can be customised by extending the velocity objects.

The *position* terminology is used throughout the API as opposed to *location* or *pose*. Pose is a term that is often used when defining a position and orientation in a three-dimensional space. However, with the support of 2D positions, this term was not favourable. Location is described by the English Oxford Dictionary as “a particular place or position”. This abstraction of “place” led us to our final decision of choosing the more precise *position* terminology.

4.2.4 Data Object. A data object represents anything that is relevant to the positioning. It can be the tracked object, the tracking object or a landmark needed for the relative positioning. Each object contains the following attributes:

- **Unique identifier:** Data objects are uniquely identified, either by a supplied identifier or a random UUID. Optionally, a developer can provide a more user-friendly display name.
- **Absolute position:** Data objects store their last known absolute position. The stored position is always relative to the global reference space introduced later in Section 4.2.5. The

relevance of this last known position can be determined using its timestamp and developers can request the transformed position in their own reference space.

- **Relative positions:** These are relative positions to other reference objects. Each object can have multiple types of positions relative to different objects. This allows a data object to have a relative distance, angle and velocity to the same object.
- **Parent object:** A data object can specify its parent. This can be useful for indicating that individual sensor objects belong to the same tracked actor.

Depending on what the data object represents, it can be extended to store the information needed for its representation. In Listing 3, we create a basic data object of a user who is uniquely identified by their e-mail. During the creation of this object, we set the current position to a geographical coordinate.

```

/* @openhps/core | version 0.2.0 */
1 const object = new DataObject("mvdewynnc@vub.be");
2 object.displayName = "Maxim Van de Wynckel";
3 object.setPosition(new GeographicalPosition(50.82075, 4.39234));

```

Listing 3: Creation of a DataObject

Data objects can be created and modified without those changes being persisted in the positioning model. In order to detect persisted changes, a listener can be added to the data object service as shown in Listing 4.

```

/* @openhps/core | version 0.2.0 */
1 const service = myModel.findDataService(DataObject);
2 service.on('insert', (uid, changedObject) => {
3   if (uid === object.uid)
4     console.log(changedObject.getPosition());
5 });

```

Listing 4: Listener for data object changes

4.2.5 Reference Space. Reference spaces are data objects that represent spaces which are used for absolute positions. Using these reference spaces, absolute positions created in a different space can easily be identified and transformed to the global reference space created when building a model.

```

/* @openhps/core | version 0.2.0 */
1 const refSpace = new ReferenceSpace(model.referenceSpace)
2   .unit(LengthUnit.CENTIMETER)
3   .translation(10, 10, 0)
4   .scale(1, 1, 0)
5   .rotation(0, 0, 0, AngleUnit.RADIANS);

```

Listing 5: Creation of a ReferenceSpace

Listing 5 shows the creation of a reference space relative to the global space represented by `model.referenceSpace`. This reference space has an origin offset. Absolute positions set when providing this reference space will automatically transform to the origin of the global space.

A reference space can transform the position, velocity and orientation in the following ways:

- **Translation:** Translate the position with an origin offset.
- **Rotation:** Rotate the position, orientation and angular velocity.
- **Scale:** Scale the position and linear velocity.
- **Perspective:** Transform the (inverse) perspective of the position (e.g. the perspective of a camera).
- **Unit conversion:** Convert the unit of a position to a reference unit.

Reference spaces can be created to model different scenarios:

- **Third-party positioning systems:** Frameworks like the WebXR [22] API manage their own origin and orientation based on the underlying hardware. The output of such third-party frameworks are high-level positions that should be aligned with the other positioning methods.
- **Sensor placement:** Developers can model a reference space for sensors that have a static offset or rotation (e.g. a motion sensor that is placed upside down).
- **Calibrated reference space:** Some sensors require a calibration (either automatic or by manual user input). A goal of OpenHPS is to easily persist this type of calibration.
- **Map storage:** As a data object, a reference space can be extended to store environment map information as outlined in our functional requirements.

In Listing 6, we set the current position of a data object to (5, 5, 5) using the previously created reference space shown in Listing 5. Internally, the stored position of myObject will be the transformed position with coordinates (-5, -5, 5).

```
/* @openhps/core | version 0.2.0 */
1 myObject.setPosition(new Absolute3DPosition(5, 5, 5), refSpace);
```

Listing 6: Setting the object position in a reference space

As these spaces are data objects, they are uniquely identified and can have a parent object or space. This parent allows for abstract reference spaces such as *rooms*, *floors* and *buildings*. These types of abstractions allow us to use different positioning methods per floor that are stored in a global reference space representing a building.

4.2.6 Services. Each positioning model can have multiple services. A service can be accessed by all nodes in that model to perform certain general actions ranging from communication services that handle the data between remote nodes, to data services that store data frames, objects or other relevant information.

A data service serialises and stores information. By default, our core API offers data services for:

- **Data objects:** To store the processed objects and their last known position. This can also be used as a persistent storage for landmarks used in the positioning.
- **Node data:** Node-specific data about DataObjects can be stored. This can be useful for intermediate calculations by noise filtering algorithms or sensor fusion techniques.

- **Trajectories:** Historical position data of DataObjects. Drivers can be implemented for storing this information in specialised databases such as MobilityDB [40].

Normal services in our framework include, but are not limited to a *time service* that allows developers to synchronise the time between multiple machines, and a *worker service* that acts as a (remote) proxy for data services.

Listing 7 shows examples of how a service can be retrieved from the model. Nodes can retrieve a data service by providing either the class of an object, an object instance or the class name of the object. This allows the use of difference services for different types of DataObjects.

```
/* @openhps/core | version 0.2.0 */
1 // Finding a data service by class
2 this.model.findDataService(DataObject);
3 // Finding a data service object instance
4 this.model.findDataService(myObject);
5 // Finding a data service by name
6 this.model.findDataService("RFDataObject");
```

Listing 7: Retrieving a data service from a model

4.2.7 Measurement Units. Unlike many positioning frameworks aiming for geographical positioning, OpenHPS aims to support a wide range of use cases ranging from small scale to celestial positioning. We provide a unit system consisting of the Unit and DerivedUnit objects. A derived unit consists of multiple units with a specific power and offset. Math.js [10] offers a similar unit system with the possibility to automatically evaluate and convert units. While this allows for the easy creation of derived units, it is not necessary for our framework.

```
/* @openhps/core | version 0.2.0 */
1 // Time unit called 'second'
2 const second = new TimeUnit('second', {
3   // Unit for 'time'
4   baseName: 'time',
5   // Also called 's', 'sec' or plural
6   aliases: ['s', 'sec', 'seconds'],
7   // Supports decimal prefixes (milli, micro, ...)
8   prefixes: 'decimal',
9 });
10
11 // Millisecond is a second with the prefix specifier milli
12 const millisecond = second.specifier(UnitPrefix.MILLI);
13
14 const minute = new TimeUnit('minute', {
15   baseName: 'time',
16   aliases: ['m', 'min', 'minutes'],
17   // Minute can be defined as 60 * 1 second
18   definitions: [{ magnitude: 60, unit: 's' }],
19 });
```

Listing 8: Unit creation

Listing 8 shows the creation of a base unit second for time. During its creation the developer can specify aliases for the unit and similar to Math.js, a unit can have a set of unit prefixes. This allows the use of “millisecond, microsecond, nanosecond, ...” without

specifically creating individual units for these specifiers. Note that aliases can be provided to optionally allow the units to be converted to string evaluators of other mathematical modules.

When creating a new unit, the developer should specify the base unit. For the minute example in Listing 8 this is done by creating a definition for converting minutes to seconds (using a magnitude of 60 for the unit seconds).

```

1  /* @openhps/core | version 0.2.0 */
2  const radSecond = new DerivedUnit('radian per second', {
3    baseName: 'angularvelocity',
4    aliases: ['rad/s', 'radians per second'],
5  })
6  .addUnit(AngleUnit.RADIAN, 1)
7  .addUnit(TimeUnit.SECOND, -1);
8
9  const degreeSecond = radSecond.swap(
10 [AngleUnit.DEGREE],
11 {
12   baseName: 'angularvelocity',
13   name: 'degree per second',
14   aliases: ['deg/s', 'degrees per second'],
15 });
16
17 const degreeMinute = radSecond.swap(
18 [AngleUnit.DEGREE, TimeUnit.MINUTE],
19 {
20   baseName: 'angularvelocity',
21   name: 'degree per minute',
22   aliases: ['deg/min', 'degrees per minute'],
23 });

```

Listing 9: Derived unit creation

In order to use a unit that is derived from other base units, a `DerivedUnit` can be created as shown in Listing 9. The developer provides a name of the unit and adds the units that are contained in the derived unit (lines 5 and 6) along with their magnitude. Variants on derived units can be created by *swapping* a unit (lines 9 and 17).

4.3 Framework Implementation

OpenHPS is implemented in TypeScript², a type-safe superset of JavaScript. It can be executed as a client-side browser application, hybrid mobile applications, on JavaScript supported embedded systems such as Espruino and even as a server-side application using Node.js³ or Deno⁴.

The ability to run our positioning model on a large range of server and client devices enables the decentralisation mentioned in the functional requirements. Additional remote components such as the socket API outlined in Section 4.3.4 allow for other programming languages to be supported as well.

4.3.1 Serialisation. Data frames and contained objects are serialisable throughout the framework. This functionality is implemented using an extension of TypedJSON⁵ that adds the ability for polymorphic data types. The detection of such data types is necessary for allowing developers to create additional position or data objects without having to recreate all classes where these are used.

²<https://www.typescriptlang.org>

³<https://nodejs.org/en/about/>

⁴<https://deno.land>

⁵<https://github.com/JohnWeisz/TypedJSON>

```

1  /* @openhps/core | version 0.2.0 */
2  {
3    "createdTimestamp": 1606501972983302,
4    "uid": "8865727c-7c98-4a8d-a33c-506d2650e59d",
5    "position": {
6      "x": -4.07093248547983,
7      "y": 55.59130128032057,
8      "timestamp": 1606502001594449,
9      "velocity": {
10     "linear": {
11       "x": -0.27608249684331726,
12       "y": 0.3606549076013354,
13       "z": 0.013291033512841348
14     },
15     "angular": {
16       "x": -3.9937982517329886,
17       "y": 0.2311694373502423,
18       "z": -0.5070813464456928
19     }
20   },
21   "orientation": {
22     "x": -0.09754179767548248,
23     "y": 0.15388368786071302,
24     "z": 0.04266920115206052,
25     "w": 0.9823363719162936
26   },
27   "unit": {
28     "name": "centimeter"
29   },
30   "referenceSpaceUID": "5582d63d-c7af-4624-9fed-6ce0d9036f62",
31   "accuracyUnit": {
32     "name": "meter"
33   },
34   "__type": "Absolute2DPosition"
35 },
36 "relativePositions": [],
37 "__type": "DataObject"
38 }

```

Listing 10: Serialised DataObject

Listing 10 shows a serialised data object created with sensor data retrieved from a *Sphero Mini*⁶ toy. The main `DataObject` and `Absolute2DPosition` have a `__type` key that defines the object type. Definitions of a unit are not included in the serialisation and its complete name is used to indicate the unit. This means that a custom unit should be available in all processes that are required to deserialise the unit.

4.3.2 Performance. One of the non-functional requirements mentioned in Section 4.1 is the ability to perform real-time data processing. In order to achieve these performance requirements, parts of the processing network can be run in their own thread, web worker or process. This threading is made possible due to the serialisability of data frames and objects, which allows the transmission of frames from one process or thread to another.

Listing 11 shows the creation of a model with parts of the graph going through a `WorkerNode`. This threaded node is initialised with a model builder function evaluated on the threaded process. If no data services are (re)initialised in this function, the data services of the main thread are made accessible in the individual threads.

A `WorkerNode` can also run a larger portion of a process network that is declared in a separate file. This is more developer friendly

⁶<https://sphero.com/products/sphero-mini>

```

1  /* @openhps/core | version 0.2.0 */
2  ModelBuilder.create()
3  .addService(/* ... */)
4  .from(/* ... */)
5  .via(new WorkerNode((builder: GraphShapeBuilder) => {
6    const { TrilaterationNode } = require('@openhps/core');
7    builder.via(new TrilaterationNode())
8  }, {
9    poolSize: 4
10  }));
11 .build().then(model => { /* ... */ });

```

Listing 11: Threaded node creation

than having to import all the nodes within a builder function. Listing 12 shows the worker node named “video” being created in the main thread (lines 2 to 6). Internally, this node is a graph created in `video.ts`. Pull requests to this node (line 8) will be forwarded to a pool of four workers.

```

1  /* @openhps/core | version 0.2.0 */
2  // main.ts //
3  ModelBuilder.create()
4  .addNode(new WorkerNode("video.ts", {
5    poolSize: 4,
6    name: "video"
7  }));
8  .from("video")
9  .via(new TimedPull(1, TimeUnit.MILLISECOND))
10 .to(/* ... */);
11 .build().then(model => { /* ... */ });
12 // video.ts //
13 export default GraphBuilder.create()
14 .from(/* ... */)
15 .via(/* ... */)
16 .to();

```

Listing 12: Threaded graph creation

As a simple demonstration of our worker node, we created a processing node calculating 5000 prime numbers for every received frame. This test was conducted on an Intel i7-6700HQ laptop CPU with 8 logical cores, running Node.js 14.10. These 5000 prime numbers can be calculated 237.03 times per second without the overhead of data frames, objects and services. The data frames that we push contain a source object, position and velocity to simulate the amount of data normally serialised and communicated between the main process and workers. However, the contained data does not affect the time it takes to compute the prime numbers.

Table 1 shows the results of our benchmark with one worker assigned to each logical CPU core. Performance is measured in frames per second (FPS) represented by the amount of computed data frames received by the sink of our model. For each worker we indicate the speed-up compared to the sequential implementation. The overhead shown with a single worker is due to the serialisation and deserialisation of data, an operation that is not required when pushing in a sequential network.

4.3.3 Precision. Calculations within the framework are made using JavaScript number operations. Time-critical operations use a

#workers	FPS	Error	Speed-up
Sequential	229.04	1.19%	-
1	200.74	0.67%	0.88
2	389.44	0.56%	1.70
3	512.42	0.92%	2.24
4	616.29	1.15%	2.69
5	671.00	0.59%	2.93
6	746.07	0.67%	3.26
7	801.32	0.90%	3.50
8	822.47	0.69%	3.59

Table 1: WorkerNode benchmark

time service that returns the time in a specific unit. This allows developers to extend the framework with additional modules such as `microtime`⁷ for more precise calculations. In addition, `Decimal.js`⁸ could be used with an extended position class to provide more precise number operations.

4.3.4 Modularity. OpenHPS provides a modular API that splits the functionality of positioning methods and algorithms in different npm⁹ modules. Using this method, developers can extend on our core framework or other components. It also prevents them from having to depend on very large modules, reducing the overall dependency size.

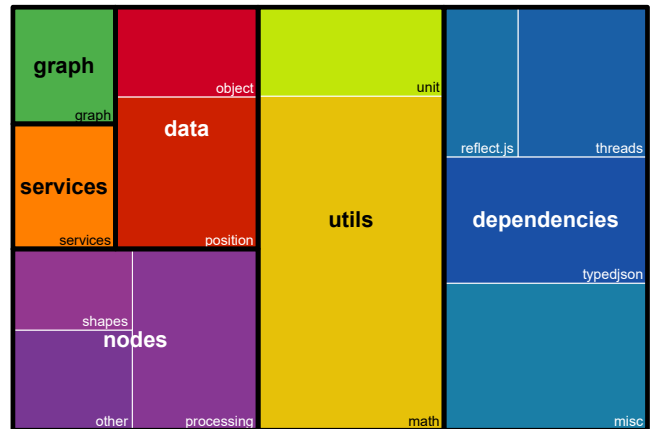


Figure 7: @openhps/core minified web module tree map

Our core API, named `@openhps/core` is available for server and web deployment in the CommonJS (CJS), ECMAScript (ESM) and Universal Module Definition (UMD) formats. Figure 7 presents an overview of the core components content size in its current state (version 0.2.0). Most of the file size is taken up by dependencies ($\approx 30\%$, indicated in blue) and the mathematical classes of `Three.js`¹⁰ ($\approx 22\%$, as part of the mathematical utilities in yellow). Default nodes such as processing nodes, graph shapes and common sink or source nodes account for $\approx 16\%$. The main purpose of the dependencies are to help with serialisation (i.e. `TypedJSON`, `Reflect.js`). Mathematical

⁷<https://www.npmjs.com/package/microtime>

⁸<https://www.npmjs.com/package/decimal>

⁹<https://www.npmjs.com>

¹⁰<https://threejs.org/docs/>

classes such as quaternions, matrices and vectors from Three.js offer general operations for handling 2D and 3D position manipulation. We list several examples of modules that can be used to extend the core functionality:

- **Data storage:** By default, the core API provides the possibility of using in-memory data storage. In order for this data to be persisted, additional components making use of different database management systems such as MongoDB, Redis or MobilityDB [40] can be applied.
- **Remote communication:** The remote APIs introduce a RemoteNode that can be added to the model. This node will transmit push (or pull) requests to a remotely connected model through either a REST API, socket connection or a message broker such as MQTT [21].
- **Positioning methods and algorithms:** The core API offers basic processing nodes to determine a position (i.e. trilateration, triangulation or fingerprinting) and can be extended with different components. Examples include techniques that require additional machine learning or computer vision libraries.
- **Symbolic positions:** Our core API offers the 2D, 3D and geographical positioning. This can be abstracted to *locations* or *places* such as a room, building or site.
- **Third-party positioning systems:** Third-party positioning solutions can be integrated into OpenHPS by using modules that provide this interface.

4.4 Demonstrator

Unlike some of the frameworks discussed in Section 3 that are made to tackle a certain issue or goal, the core idea of OpenHPS is to combine different positioning concepts into one model.

As a non-trivial demonstrator, we provide a positioning system for a Sphero Mini toy using the internal sensors and an external Logitech Brio¹¹ camera. The Sphero provides raw sensor reading for the linear and angular velocity, raw accelerometer data, orientation and an internally computed position. This internal position is computed by the Sphero toy itself and makes use of the motor velocity, accelerometer and gyroscope.

We make use of the @openhps/core¹², @openhps/opencv¹³ and the use case-specific @openhps/sphero¹⁴ modules to construct a model that fuses these multiple sources together into one position. The model consists of four sources; the video input, internally computed position, the input that is sent to the Sphero and finally the dead reckoned position that is calculated by the framework itself using the provided velocity.

Our setup is shown in Figure 8. We used yellow floor markers to define an area of 260 cm × 200 cm. The camera is positioned with a perspective view on the area and the start position of the Sphero is at the bottom right corner of the camera source.

For the scope of this demonstration, the Sphero performs a simple trajectory. The input for this device consists of an orientation (heading) and speed. Before the start of our input trajectory, we



Figure 8: Demonstrator overview

manually calibrated the origin orientation using the provided mobile application for the Sphero. This provides us knowledge on the start orientation used by the internally calculated position which allowed us to define the reference spaces.

Various methods exist to combine the aforementioned positioning methods. Figure 9 shows the simplified graph presentation of our demonstrated positioning model. Starting from the four different sources, we will discuss how each signal is processed and fused together. We use two feedback loops from our fused position to provide temporal information to our positioning model.

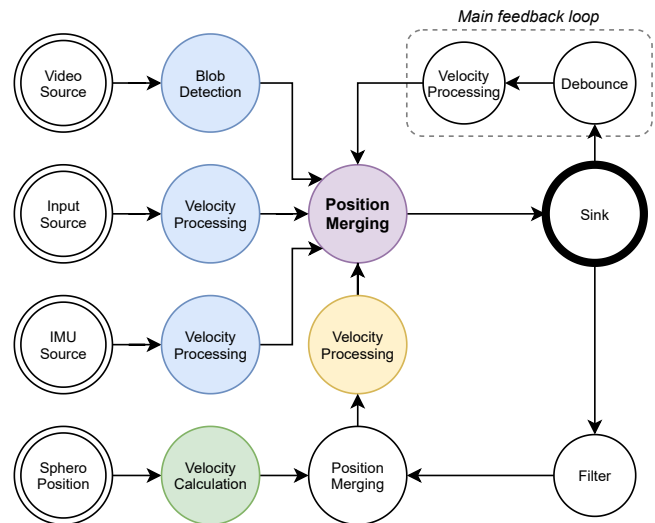


Figure 9: Demonstrator positioning model

The results of each independent source is shown in the trajectory scatter plots in Figure 10. Each positioning method has a different frequency, resulting in a varying amount of data points used to determine the position. Our main feedback loop in Figure 9 ensures that the fused position never relies on a single source.

¹¹<https://www.logitech.com/en-us/product/brio>

¹²<https://github.com/OpenHPS/openhps-core/>

¹³<https://github.com/OpenHPS/openhps-opencv/>

¹⁴<https://github.com/OpenHPS/openhps-sphero/>

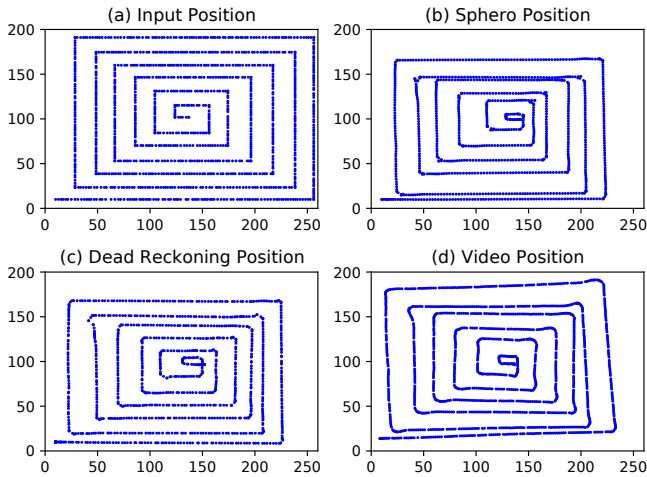


Figure 10: Individual position estimates for the given input (a), including the internally calculated position (b), dead reckoning position (c) and video source (d)

4.4.1 Input Control. Input to the Sphero is given using a heading (degrees), speed (0-255) and roll duration. We make the assumption that the Sphero Mini has a maximum speed of 1 m/s as documented on the product website¹⁵.

As input trajectory, we provide a spiralling rectangle starting from an outer corner to the centre of the area with a speed of $150 (= 0.58\text{ m/s})$. We provide a basic roll duration of $4.2\text{ seconds} (= 2.436\text{ m})$ for the X-axis and a roll duration of $3.2\text{ seconds} (= 1.856\text{ m})$ for the Y-axis. Every turn, the duration of the movement along the X-axis is reduced by 168 ms while the movement along the Y-axis is reduced by 128 ms . This input is fed to our framework’s velocity processing node resulting in the output shown in Figure 10a.

4.4.2 Visual Positioning. The video source uses the OpenCV [5] library to capture a 30 FPS camera feed from the Logitech Brio camera which has a perspective view of the floor. When processing the video stream, we create the inverse perspective view by manually specifying the position of four yellow markers on the floor. This creates a wrapped image rectangle of $1040\text{ px} \times 800\text{ px}$.

Once the video source is wrapped, blob detection is used to determine the centroid position of the blue Sphero Mini. We apply a colour mask that converts the image to an HSV colour space and performs a masking filter to only show the blue ball as illustrated in Figure 11. Next, in Listing 13 we create a custom processing node that sets the position of our tracked object to the pixel position of the blob. As the accuracy for our position we take the square root of the blob area. A reference space is created (lines 1 to 4 in Listing 14) and applied to the output position (pixel coordinate) on line 25 to scale it to the corresponding rectangle dimensions.

Without interference from other sources, the video processing provides the output shown in Figure 10d. We will use this source as our most accurate position, as it is the only available external positioning method.

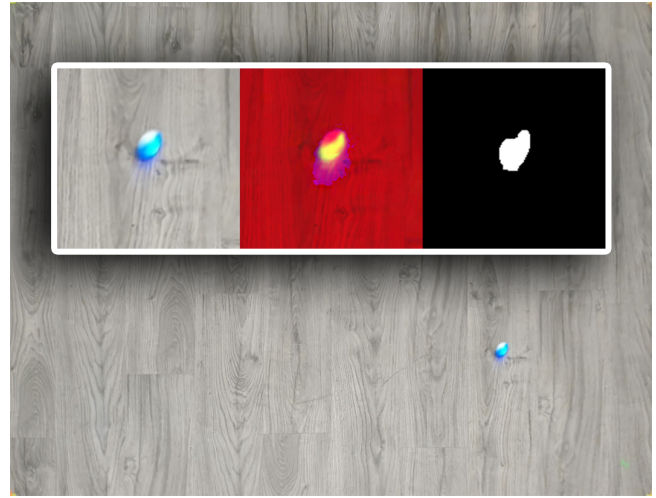


Figure 11: Conversion of wrapped video to blob

4.4.3 Internal Position. In Figure 10b we show the internal positioning calculated by the Sphero, converted to a certain reference space created with our calibrated orientation knowledge. Instead of using the raw position, we determine the displacement of this internal position (using the filtered feedback loop shown in Figure 9) and apply this displacement to the fused position.

```

1  /* @openhps/core | version 0.2.0 */
2  class ContourDetectionNode extends ProcessingNode<VideoFrame> {
3  public process(frame: VideoFrame): Promise<VideoFrame> {
4  return new Promise((resolve) => {
5  let contours = frame.image.findContours(
6  OpenCV.RETR_EXTERNAL,
7  OpenCV.CHAIN_APPROX_SIMPLE);
8  if (contours.length >= 1) {
9  // Sort contours by area
10 contours = contours.sort((a, b) => a.area - b.area);
11 // Select the contour with the largest area size
12 const m = contours[0].moments();
13 const center = new OpenCV.Vec2(
14 m.m10 / m.m00,
15 m.m01 / m.m00);
16 // Use the center as the 2D pixel position
17 const position = new Absolute2DPosition(
18 center.x,
19 center.y);
20 position.unit = LengthUnit.CENTIMETER;
21 position.accuracy = Math.sqrt(contours[0].area);
22 frame.source.setPosition(position);
23 }
24 resolve(frame);
25 });
26 }

```

Listing 13: Contour detection processing node

4.4.4 Dead Reckoning Position. Apart from an internally calculated position, the Sphero provides raw sensor data for the accelerometer, gyroscope, orientation and velocity (internally fused from the motor velocity and acceleration). For the scope of this demonstration we make use of this velocity and orientation to compute the position using OpenHPS. The output of this source is shown in Figure 10c.

¹⁵<https://support.sphero.com/article/6drb2qggx4-sphero-mini-faq>

```

1  /* @openhps/core | version 0.2.0 */
2  const videoSpace = new ReferenceSpace(defaultSpace)
3    .translation(1040, 800)
4    .rotation(new Euler(180, 180, 0, 'ZXY', AngleUnit.DEGREE))
5    .scale(4, 4);
6  /* ... */
7  export default GraphBuilder.create()
8    .from(new VideoSource(new CameraObject("sphero_video"), {
9      autoPlay: true,
10     fps: 30,
11     // Do not fetch a frame if the webcam can not handle it
12     throttleRead: true,
13     source: new CameraObject("sphero_video")
14   })).load("/dev/video2"))
15   .via(new ImageTransformNode({
16     src: [
17       new OpenCV.Point2(307, 120),
18       new OpenCV.Point2(1473, 87),
19       new OpenCV.Point2(1899, 891),
20       new OpenCV.Point2(20, 1024),
21     ],
22     height: 800, // 200cm
23     width: 1040 // 260cm
24   }))
25   .via(new ColorMaskProcessing({
26     minRange: [90, 50, 50],
27     maxRange: [140, 255, 255]
28   }))
29   .via(new ContourDetectionNode())
30   .convertFromSpace(videoSpace)
31   .to();

```

Listing 14: Graph shape video

4.4.5 Model Creation. In Listing 15 we combine the four graph shapes for our video output, internal position, input and dead reckoned position. We use a built-in object merging node (lines 22 to 30) that merges frames where the source UID is equal to “sphero”. The merge node will wait until all of its incoming edges pushed a frame, or the timeout of 20 ms has been reached. By default, this merge will use the weighted average of all incoming positions, velocities and orientations (with the weight being the inverse of its accuracy). Developers have the choice to choose their own strategy by, for instance, selecting a single position based on the highest accuracy.

This final fused position is presented in Figure 12a. Compared to the individual positioning methods shown in Figure 10, we have more data points for our positions. This is because we do not wait for all sources to provide data before computing the next position (20 ms timeout). Our feedback loop called “feedback” ensures that position fusion never relies on just one source.

4.4.6 Evaluation. We have shown our completed positioning system in the previous section. Four sources and a feedback loop resulted in a fused position. In order to evaluate this positioning model, we removed parts of our video source to simulate an obstacle or blind spots for the camera.

The goal of this evaluation is to first ensure that the positioning model can function with missing information and to determine the error as a result of this missing positioning data.

To illustrate a baseline of the remaining sources that will take over the positioning, we show the merged position of all sources except the video source in Figure 12b.

Figures 12c and 12d show two examples with video blind spots (grey areas). Indicated in blue are the data points where the video

```

1  /* @openhps/core | version 0.2.0 */
2  ModelBuilder.create()
3    .addNode(new WorkerNode("video.ts", {
4      poolSize: 1,
5      name: "video"
6    })))
7    .addShape(inputSource)
8    .addShape(spheroPosition)
9    .addShape(spheroVelocity)
10   // Feedback loop
11   .addShape(GraphBuilder.create()
12     .from("merged")
13     .debounce(10, TimeUnit.MILLISECOND)
14     // Clone the frame and update timestamp
15     // (needed to process velocity)
16     .clone({
17       repack: true
18     })
19     .via(new VelocityProcessingNode())
20     .to("feedback"))
21   .addShape(GraphBuilder.create()
22     .from("video", "sphero_position", "input",
23       "sphero_velocity", "feedback")
24     .merge((frame, options) => options.sourceNode,
25       {
26         timeout: 20,
27         timeoutUnit: TimeUnit.MILLISECOND,
28         // Minimum two sources, else the feedback
29         // loop will continue
30         minCount: 2,
31         objectFilter: obj => obj.uid === 'sphero',
32       })
33     .via("merged") // Feedback loop
34     .to(new CSVDataSink("position.csv", [
35       { id: "timestamp", title: "timestamp" },
36       { id: "x", title: "x" },
37       { id: "y", title: "y" },
38     ]), (frame: DataFrame) => {
39       return {
40         timestamp: frame.createdTimestamp,
41         x: frame.source.getPosition().toVector3().x,
42         y: frame.source.getPosition().toVector3().y,
43       });
44     })))
45   .build().then(model => {
46     // Model created
47   });

```

Listing 15: Demonstration model creation

processing was still able to detect an object, whereas the positions calculated without input from the video source are highlighted in red in the figures.

Source(s)	Avg error	Max error
all sources (Fig. 12a)	0.00 cm	0.00 cm
input control only (Fig. 10a)	23.07 cm	50.06 cm
internal position only (Fig. 10b)	16.16 cm	33.38 cm
dead reckoning only (Fig. 10c)	17.09 cm	34.44 cm
video source only (Fig. 10d)	1.30 cm	4.74 cm
all sources excl. video (Fig. 12b)	13.59 cm	29.73 cm
blind spot left (Fig. 12c)	4.26 cm	21.65 cm
blind spot right (Fig. 12d)	4.81 cm	24.40 cm

Table 2: Average and maximum XY position error compared to the fused position with all sources

In Table 2, we show the average and maximum position error compared to the final fused position from Figure 12a. This error is

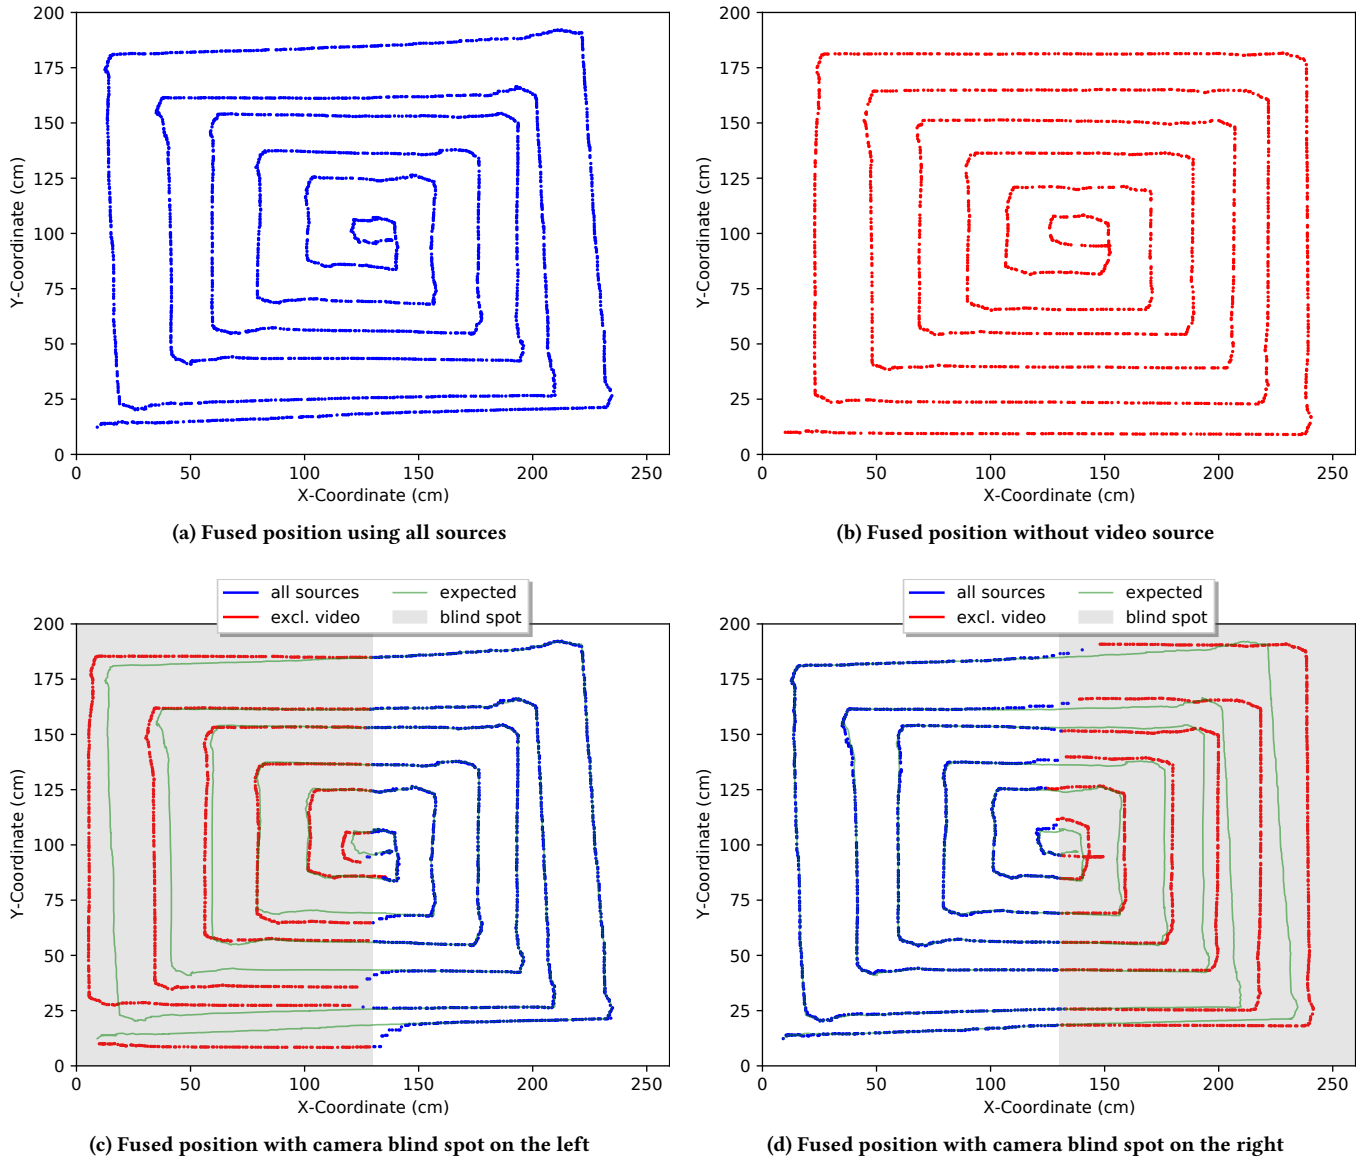


Figure 12: Fused positions processed by our model

determined by taking 100 timestamped key points in each trajectory (every 51 ms) and calculating the average and maximum difference for those points.

Our results show that the video source is the main positioning method in the fused position. Blind spots in this source result in the model falling back to the remaining dead reckoning. However, the positioning model is self correcting and will gradually align with the video source position once it becomes available.

The positioning model we illustrated in Figure 9 is highly adaptable depending on the desired outcome. For example, noise filtering nodes such as a Simple Moving Average (SMA) can be used on the video accuracy to provide a smoother transition at the border of the blind spot.

With the evaluation in Figure 12 and Table 2 we have proven that multiple producers of sensor information can be merged together into a continuous stream of fused positions. By creating blind spots in our video source, we have shown that the model is capable of running without our main visual positioning method.

5 CONCLUSION AND FUTURE WORK

We have presented OpenHPS, an open source hybrid positioning system. We focused on the different actors of our system that have been defined based on an investigation of some of the more prominent existing positioning methods and algorithms. These actors, in combination with our requirements, were used in developing

our positioning framework with its graph topology. We further presented our definition of *nodes*, *data frames*, *data objects* as well as *positions*. Finally, the OpenHPS implementation in TypeScript highlighting how we addressed and satisfied our non-functional requirements, has been discussed in Section 4.3.

In the demonstrator application in Section 4.4, we have illustrated how multiple positioning methods can be fused via some high-level decision fusion. We have further highlighted—by removing certain parts of our main sensor source—how the presented positioning model continues to work on the remaining positioning methods and manages to recover once the input from the main sensor source is back.

A major effort in the design and development of OpenHPS went into the extensibility of our framework. External modules can be used to extend OpenHPS with additional positioning methods and techniques. Some basic positioning methods are currently included in the core OpenHPS component. However, in order to prevent that the core contains potentially unused nodes, in the future some of these basic positioning methods and algorithms might be moved to their own dedicated modules (e.g. for fingerprinting techniques). Apart from individual nodes, these modules can also provide complete graph shapes that act similar to position providers in other high-level hybrid positioning systems.

The real-time processing of positioning information was the most important goal for the presented OpenHPS framework. During the development, the computing performance of the positioning model has therefore always received a high priority and lead to the introduction of worker nodes and services. Future research and development of the OpenHPS hybrid positioning system might focus on optimising the serialisation of data frames in order to only serialise changes in data rather than all available information. This optimisation would ensure that data transfers are limited to new data only.

Overall, the presented OpenHPS framework represents a solid hybrid positioning solution offering various possibilities for future extensions. An obvious future extension would be the introduction of additional layers of abstraction providing similar high-level functionality as offered by some of the solutions discussed in the related work. Further, the exiting reference spaces might be extended in a separate OpenHPS module in order to represent and deal with symbolic locations, similar as offered by HyLocSys [15]. The support of symbolic locations [19] will further strengthen the position of OpenHPS as a framework for context-aware computing and implicit human-computer interaction.

REFERENCES

- [1] Martin Abadi et al. 2016. TensorFlow: A System for Large-scale Machine Learning. In *Proceedings of USENIX 2016, International Conference on Operating Systems Design and Implementation*. Savannah, USA.
- [2] Sean J. Barbeau, Miguel A. Labrador, Philip L. Winters, Rafael Pérez, and Nevine Labib Georggi. 2008. Location API 2.0 for J2ME: A New Standard in Location for Java-enabled Mobile Phones. *Computer Communications* 31, 6 (April 2008). <https://doi.org/10.1016/j.comcom.2008.01.045>
- [3] Stephane Beaugard and Harald Haas. 2006. Pedestrian Dead Reckoning: A Basis for Personal Positioning. In *Proceedings of WPNC 2006, Workshop on Positioning, Navigation and Communication*. Merida City, Mexico. <https://doi.org/10.1109/ICEE.2011.6106608>
- [4] Anja Bekkelien and Michel Deriaz. 2012. Hybrid Positioning Framework for Mobile Devices. In *Proceedings of UPINLBS 2012, International Conference on Ubiquitous Positioning, Indoor Navigation, and Location Based Service*. Helsinki, Finland. <https://doi.org/10.1109/UPINLBS.2012.6409759>
- [5] Gary Bradski. 2000. The OpenCV Library. *Dr. Dobbs's Journal of Software Tools* (November 2000). <https://www.drdoobs.com/open-source/the-opencv-library/184404319#>
- [6] Nirupama Bulusu, John Heidemann, and Deborah Estrin. 2000. GPS-less Low-Cost Outdoor Localization for Very Small Devices. *IEEE Personal Communications* 7, 5 (October 2000). <https://doi.org/10.1109/98.878533>
- [7] Harry Chen, Tim Finin, and Anupam Joshi. 2003. An Ontology for Context-aware Pervasive Computing Environments. *The Knowledge Engineering Review* 18, 3 (September 2003). <https://doi.org/10.1017/S0269888904000025>
- [8] Zhenghua Chen, Han Zou, Hao Jiang, Qingchang Zhu, Yeng Soh, and Lihua Xie. 2015. Fusion of WiFi, Smartphone Sensors and Landmarks Using the Kalman Filter for Indoor Localization. *Sensors* 15, 1 (January 2015). <https://doi.org/10.3390/s150100715>
- [9] Adam L. Davis. 2019. Akka Streams. In *Reactive Streams in Java*. https://doi.org/10.1007/978-1-4842-4176-9_6
- [10] Mansfield E. de Jong J. 2014. Math.js: An Advanced Mathematics Library for JavaScript. *Computing in Science & Engineering* 20, 1 (January 2014). <https://doi.org/10.1109/MCSE.2018.01111122>
- [11] B. Louis Decker. 1986. World Geodetic System 1984. In *Proceedings to the Fourth International Geodetic Symposium on Satellite Positioning*. Austin, USA.
- [12] Cristiano di Flora, Massimo Ficco, Stefano Russo, and Vincenzo Vecchio. 2005. Indoor and Outdoor Location Based Services for Portable Wireless Devices. In *International Workshop on Services and Infrastructure for Ubiquitous and Mobile Internet*. Columbus, USA. <https://doi.org/10.1109/ICDCSW.2005.77>
- [13] Goran M. Djuknic and Robert E. Richton. 2001. Geolocation and Assisted GPS. *Computer* 34, 2 (February 2001). <https://doi.org/10.1109/2.901174>
- [14] Wilfried Elmenreich. 2002. *An Introduction to Sensor Fusion*. Technical Report 47/2001. Vienna University of Technology.
- [15] Massimo Ficco and Stefano Russo. 2009. A Hybrid Positioning System for Technology-independent Location-aware Computing. *Software: Practice and Experience* 39, 13 (September 2009). <https://doi.org/10.1002/spe.919>
- [16] Marc Geilen and Twan Basten. 2004. Reactive Process Networks. In *Proceedings of EMSOFT 2004, International Conference on Embedded Software*. Pisa, Italy. <https://doi.org/10.1145/1017753.1017778>
- [17] Yanying Gu, Anthony Lo, and Ignas Niemegeers. 2009. A Survey of Indoor Positioning Systems for Wireless Personal Networks. *IEEE Communications Surveys & Tutorials* 11, 1 (2009). <https://doi.org/10.1109/SURV.2009.090103>
- [18] Janne Haverinen. 2014. Utilizing Magnetic Field Based Navigation. US Patent 8,798,924.
- [19] Jeffrey Hightower and Gaetano Borriello. 2001. Location Systems for Ubiquitous Computing. *Computer* 34, 8 (August 2001). <https://doi.org/10.1109/2.940014>
- [20] Fabian Hölzke, Johann-P. Wolff, and Christian Haubelt. 2019. Improving Pedestrian Dead Reckoning Using Likely Paths and Backtracking for Mobile Devices. In *Proceedings of PerLS 2019, International Workshop on Pervasive Smart Living Spaces*. Kyoto, Japan. <https://doi.org/10.1109/PERCOMW.2019.8730734>
- [21] Urs Hunkeler, Hong Linh Truong, and Andy Stanford-Clark. 2008. MQTT-S: A Publish/Subscribe Protocol for Wireless Sensor Networks. In *Proceedings of the International ICST Workshop on Intelligent Networks: Adaptation, Communication & Reconfiguration*. Bangalore, India. <https://doi.org/10.1109/COMSWA.2008.4554519>
- [22] Brandon Jones and Nell Waliczek. 2020. WebXR Device API. <https://www.w3.org/TR/webxr/>
- [23] John Krumm, Steve Harris, Brian Meyers, Barry Brumitt, Michael Hale, and Steve Shafer. 2000. Multi-Camera Multi-Person Tracking for EasyLiving. In *Proceedings of VS 2000, International Workshop on Visual Surveillance*. Dublin, Ireland. <https://doi.org/10.1109/VIS.2000.856852>
- [24] Axel Küpper. 2005. *Location-based Services: Fundamentals and Operation*. John Wiley & Sons.
- [25] Edward A. Lee and Thomas M. Parks. 1995. Dataflow Process Networks. *Proc. IEEE* 83, 5 (May 1995), 773–801. <https://doi.org/10.1109/5.381846>
- [26] Hui Liu, Houshang Darabi, Pat Banerjee, and Jing Liu. 2007. Survey of Wireless Indoor Positioning Techniques and Systems. *IEEE Transactions on Systems, Man, and Cybernetics* 37, 6 (2007). <https://doi.org/10.1109/TSMCC.2007.905750>
- [27] Blair MacIntyre and Trevor F. Smith. 2018. Thoughts on the Future of WebXR and the Immersive Web. In *Proceedings of International Workshop on Creativity in Design with & for Mixed Reality*. Munich, Germany. <https://doi.org/10.1109/ISMAR-Adjunct.2018.00099>
- [28] Ellon Mendes, Pierrick Koch, and Simon Lacroix. 2016. ICP-based Pose-Graph SLAM. In *Proceedings of SSRR 2016, International Symposium on Safety, Security, and Rescue Robotics*. Lausanne, Switzerland. <https://doi.org/10.1109/SSRR.2016.7784298>
- [29] Piotr Mirowski, Tin Kam Ho, Saehoon Yi, and Michael MacDonald. 2013. SignalSLAM: Simultaneous Localization and Mapping with Mixed WiFi, Bluetooth, LTE and Magnetic Signals. In *Proceedings of IPIN 2013, International Conference on Indoor Positioning and Indoor Navigation*. Montbeliard-Belfort, France. <https://doi.org/10.1109/IPIN.2013.6817853>
- [30] Sudeep Pasricha, Viney Ugave, Charles W Anderson, and Qi Han. 2015. LearnLoc: A Framework for Smart Indoor Localization with Embedded Mobile Devices.

- In *Proceedings of CODES 2015, International Conference on Hardware/Software Codesign and System Synthesis*. Amsterdam, Netherlands. <https://doi.org/10.1109/CODESIS.2015.7331366>
- [31] Andrei Popescu. 2016. Geolocation API Specification 2nd Edition. <https://www.w3.org/TR/geolocation-API/>
- [32] Morgan Quigley, Ken Conley, Brian Gerkey, Josh Faust, Tully Foote, Jeremy Leibs, Rob Wheeler, and Andrew Y Ng. 2009. ROS: An Open-Source Robot Operating System. In *Proceedings of the International Workshop on Open Source Software*. Kobe, Japan.
- [33] Shahram Rezaei and Raja Sengupta. 2007. Kalman Filter-based Integration of DGPS and Vehicle Sensors for Localization. *IEEE Transactions on Control Systems Technology* 15, 6 (October 2007). <https://doi.org/10.1109/TCST.2006.886439>
- [34] Wilson Sakpere, Michael Adeyeye-Oshin, and Nhlanhla B.W. Mlitwa. 2017. A State-of-the-Art Survey of Indoor Positioning and Navigation Systems and Technologies. *South African Computer Journal* 29, 3 (December 2017). <https://doi.org/10.18489/sacj.v29i3.452>
- [35] Albrecht Schmidt, Michael Beigl, and Hans-Werner Gellersen. 1999. There is More to Context Than Location. *Computers & Graphics* 23, 6 (December 1999). <https://doi.org/10.1109/TCST.2006.886439>
- [36] Philipp M. Scholl, Stefan Kohlbrecher, Vinay Sachidananda, and Kristof Van Laerhoven. 2012. Fast Indoor Radio-Map Building for RSSI-based Localization Systems. In *Proceedings INSS 2012, International Conference on Networked Sensing*. Antwerp, Belgium. <https://doi.org/10.1109/INSS.2012.6240574>
- [37] Stefan Steiniger, Moritz Neun, and Alistair Edwardes. 2006. Foundations of Location Based Services.
- [38] Shinya Sumikura, Mikiya Shibuya, and Ken Sakurada. 2019. OpenVSLAM: A Versatile Visual SLAM Framework. In *Proceedings of MM 2019, International Conference on Multimedia*. Nice, France. <https://doi.org/10.1145/3343031.3350539>
- [39] Wendong Xiao, Wei Ni, and Yue Khing Toh. 2011. Integrated Wi-Fi Fingerprinting and Inertial Sensing for Indoor Positioning. In *Proceedings of IPIN 2011, International Conference on Indoor Positioning and Indoor Navigation*. Guimaraes, Portugal. <https://doi.org/10.1109/IPIN.2011.6071921>
- [40] Esteban Zimányi, Mahmoud Sakr, Arthur Lesuisse, and Mohamed Bakli. 2019. MobilityDB: A Mainstream Moving Object Database System. In *Proceedings of SSTD 2019, International Symposium on Spatial and Temporal Databases*. Vienna, Austria. <https://doi.org/10.1145/3340964.3340991>