Computer Networks 88 (2015) 202–217

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

Emitter localisation from reception timestamps in asynchronous networks

Nicolò Facchi[a], Francesco Gringoli[a], Fabio Ricciato[b,c,*], Andrea Toma[c]

[a] Department of Information Engineering, CNIT—University of Brescia, Brescia, Italy
[b] Faculty of Computer and Information Science, University of Ljubljana, Ljubljana, Slovenia
[c] Austrian Institute of Technology, Vienna, Austria

ARTICLE INFO

Article history:
Received 4 November 2014
Revised 11 April 2015
Accepted 9 June 2015
Available online 16 June 2015

Keywords:
Time-based localisation
Asynchronous networks
Opportunistic localisation
WiFi testbed

ABSTRACT

We address the problem of localising a mobile terminal ("blind" node) in unknown position from a set of "anchor" nodes in known positions. The proposed method does not require any form of node synchronisation nor measurement (or control) of the transmission times, which is difficult or anyway costly to achieve in practice. It relies exclusively on reception timestamps collected by the anchor nodes, according to their local clocks, that overhear packets transmitted by the blind node and by (at least one) other transmitting node(s) in known position, e.g., other anchors. The clock differences between the nodes are not eliminated *ex ante* through clock synchronisation, as in traditional ToA and TDoA methods. Instead, they are counteracted *ex post*, during the data processing stage, leveraging the data redundancy that is intrinsic to the multiple reception of the same packet by different (anchor) nodes. We validate the proposed method in different experimental settings, indoor and outdoor, using exclusively low-cost Commercial-Off-The-Shelf WiFi devices, achieving sub-metre accuracy in full line-of-sight conditions and metre-level accuracy in mild non-line-of-sight environment. The proposed method does not require that the blind node participate actively to the localisation procedure and can use "opportunistically" any legacy signal or packet available over-the-air for communication purposes. Considering the very minimal requirement on the system—basically, only that anchors in known positions are able to collect and share reception timestamps—the proposed approach can enable practical adoption of opportunistic and/or co-operative localisation on top of existing radio communication systems.

1. Introduction

Despite the widespread success of Global Navigation Satellite Systems (GNSS), alternative radio-based positioning systems remain of interest for those applications where GNSS signals are not available (e.g., indoor, tunnels, urban canyon) and/or the integration of GNSS receivers is not feasible due to size, energy or cost constraints. In order to preserve cost-effectiveness and ease of adoption, it is desirable to develop localisation methods that reuse *opportunistically* the legacy radio signals, protocols and receivers which are anyway available for communication purposes. We consider the problem of localising a transmitting node in unknown position, called "blind node" hereafter, based on the reception measurements collected by a set of nodes in known positions, called "anchors".

One possible approach is to rely on *power* measurements, leveraging the Received Signal Strength Indicator (RSSI) available today in almost all Commercial Off-The-Shelf (COTS) transceivers to deploy RSSI-based localisation functions on top of existing wireless communication

---

[*] Corresponding author at: Faculty of Computer and Information Science, University of Ljubljana, Večna pot 113, SI-1001 Ljubljana, Slovenia. Tel.: +386 14798285.

E-mail address: fabio.ricciato@fri.uni-lj.si (F. Ricciato).

http://dx.doi.org/10.1016/j.comnet.2015.06.003
1389-1286/

systems. Both commercial and open-source tools are available for several technologies, including WiFi, Bluetooth, and IEEE 802.15.4 (e.g. [1–3]). Generally speaking, RSSI-based methods exhibit low accuracy, since in real environments the received power is heavily affected by other factors than the distance from the transmitter, most prominently multipath propagation and antenna patterns. Several previous work seek to improve accuracy of RSSI-based methods using machine learning techniques that create and use detailed power loss maps for specific environments [4–7], in the direction of so called "RSSI fingerprinting" techniques.

Another class of approaches is based on *time* measurements. As signals travel at constant speed, the relation between reception time and distance is tighter than for received power, therefore time-based localisation is intrinsically superior to RSSI-based methods in terms of potential accuracy [8]. However, the traditional approaches to time-based localisation require the system to support some form of *node synchronisation* and/or accurate *measurement or control of transmission times*. Both these functionalities are difficult or costly to achieve in practice [9]. Node synchronisation requires the distribution of synchronisation signals, either via additional wired infrastructure or wireless protocols, and the implementation of dedicated functions into the devices which consume additional energy, bandwidth, and computation resources [9]. Controlling or measuring the actual transmission time is also difficult with standard hardware, due to unpredictable delay components between the construction of the packet and the transfer to the network interface on one hand, and the intrinsic variability of channel access delays due, e.g., to MAC dynamics [9]. For these reasons we are interested in time-based localisation methods that *(i)* do not require any synchronisation across nodes and *(ii)* do not rely on knowledge nor control of *transmission* times.

In standard time-of-arrival (ToA) methods all nodes must use a common synchronous clock reference and must take active part to the localisation procedure [8,10,11]. The transmitting node attaches a transmission timestamp to the packet in order to allow the receiver to compute the travel time (also called "time of flight") and from there derive the distance from the transmitter (one-way ranging). In two-way ranging methods instead, node A sends a probe packet to node B and the latter answers back with a reply packet that includes the information about the elapsed time between the reception of the probe and the transmission of the reply packet (also called "turn-around time"). By subtracting this value from the elapsed time between the transmission of the probe packet and reception of the reply packet, node A is able to compute the distance (range) to B. Two-way ranging does not need clock synchronisation between the nodes, however it still requires accurate knowledge of (the difference between) pairs of *transmission and reception times*. Moreover, it requires that the node to be localised ("blind node") take active part into the ranging phase [12]. With time-difference of arrival (TDoA) methods an emitter node can be localised by a set of anchors in known positions that measure exclusively the reception times of the signals transmitted by the blind node [13]. This method relaxes the requirement on the transmission times, but still needs tight synchronisation between the anchor nodes for accurate localisation [14].

Due to the requirements on node synchronisation and/or measurement or control of transmission time, it is impractical to deploy these methods on COTS devices [15–18]. To overcome these limitations, we consider here a localisation method that *relies exclusively on reception timestamps and does not require any form of node synchronisation*. The proposed method can passively localise a (possibly unaware) blind node while it is executing standard communication operations, i.e., exchanging packets with a neighbouring node, without requiring its active participation to the localisation procedure. In other words, our method meets all the following system-level requirements:

- The blind node transmits beacons and/or packets (e.g. towards one or more other nodes, mobile or fixed) but it does not take any active part to the localisation process. The localisation procedure involves exclusively a set of "anchor" nodes in known positions.
- No synchronisation mechanism is in place between the anchor nodes nor between blind and anchors, i.e., we assume *asynchronous nodes*.
- Anchor nodes can measure only local *reception* times (i.e., timestamps) while the *transmission* times are unknown and cannot be controlled, i.e., we assume *asynchronous signals*.

It can be easily seen that the above requirements do not allow to adopt ToA, two-way ranging, TDoA nor other hybrid variants proposed recently in the literature (e.g. [18–24]) that anyway demand some combination of node synchronisation, knowledge and/or control of transmission times and blind node involvement.

By fulfilling the above requirements, the method described here can be adopted in a wide range of practical applications to localise passively a mobile transmitter based exclusively on the *local reception timestamps* collected by other nodes placed in the surrounding environment. Waiving the requirements of node and signal synchronisation enables the implementation of localisation capabilities in legacy communication systems, which are typically completely asynchronous. Possible applications include indoor localisation of smartphones, customer tracking for retail marketing intelligence [25], localisation in wireless sensor networks (WSN), and extension (augmentation) of GNSS positioning in Cooperative Intelligent Transportation Systems (C-ITS) [26].

While in traditional localisation systems clock differences are eliminated *ex-ante* by the synchronisation mechanism, we consider a model where clock differences are accounted for *ex-post*, i.e., they are either cancelled or estimated during the data processing stage, jointly with the estimation of the (unknown) blind position. This approach is conceptually similar to the well-known "double-difference" method used in GNSS to eliminate residual errors on the carrier *phase* for high-precision applications (see e.g. [27]). Only a few recent pioneering works have started to consider this class of techniques for *time* (instead of phase) measurements in the context of terrestrial (not GNSS) localisation systems [28–34]. The method elaborated here belongs to the class of so-called *differential* time-difference of arrival (DTDoA). A detailed review of previous DTDoA literature is given later in the Related Work section. The main goal of our contribution is to demonstrate that DTDoA localisation can deliver sub-metre