in Proc. IEEE ICASSP 2014
Florence, Italy, May 2014, pp. 1866-1870

ROBUST ASYNCHRONOUS INDOOR LOCALIZATION USING LED LIGHTING

Georg Kail[1], *Patrick Maechler*[2], *Nicholas Preyss*[1], *Andreas Burg*[1]

[1]École polytechnique fédérale de Lausanne (EPFL), Switzerland; georg.kail@epfl.ch, nicholas.preyss@epfl.ch, andreas.burg@epfl.ch
[2]Eidgenössische Technische Hochschule Zürich (ETH Zürich), Switzerland; maechler@iis.ee.ethz.ch

ABSTRACT

We propose a low-cost system for indoor self-localization of mobile devices using modulated LED ceiling lamps that are fully autonomous and broadcast their identifiers without any synchronization. The proposed self-localization method is designed to handle this lack of synchronization as well as the possibility of blocked line-of-sight connections or severe attenuation in real-world environments. This robustness is achieved by applying a suitable Bayesian signal model and by taking into account the inherent sparsity in detecting the concurrently visible lamps. The proposed estimator of the location approximates optimal Bayesian estimation while maintaining low complexity. Simulation results confirm a significant gain in performance compared to a classical matched-filter approach.

Index Terms— Indoor localization, visible light, VLC, LED, sparse signal recovery

1. INTRODUCTION

Indoor self-localization of mobile devices is a topic that has recently attracted increasing attention. Its various applications include indoor navigation as well as location-aware services and advertisements in large public buildings such as museums or shopping malls [1]. Satellite-based localization systems such as the global positioning system (GPS), which are widely used for precise outdoor localization, are often unavailable in indoor environments where signals from satellites are strongly attenuated or affected by multipath propagation. A number of alternative techniques are, thus, used for indoor localization; examples include systems based on ultra-wideband, infra red, WiFi, radio frequency identification (RFID), Bluetooth, and ultrasound [2]. This wide range of technologies varies significantly in localization precision, complexity, and price.

With the forthcoming change in lighting from traditional incandescent lamps to light-emitting diodes (LEDs), visible light is becoming an interesting option for indoor localization. Since LEDs are suitable for modulation at high rates invisible to the human eye, each lamp can be designed to act as a beacon by broadcasting a unique identifier. Recording such identifiers allows a mobile device to localize itself around the known locations of the beacons.

Visible light technology offers a number of inherent advantages compared to radiofrequency (RF) systems. One significant benefit of using light is that it generally does not penetrate walls, which allows correct detection of the room with high probability. Furthermore, light suffers less from multipath effects than RF signals, and interference from other wireless systems is low. At the same time, visible light based systems do not generate RF interference and thus avoid congesting the limited ISM frequency bands. This also makes them applicable to indoor environments where such interference is of concern, e.g., due to the presence of other sensitive electronic devices. Moreover, in places where lighting is needed for its own sake, the expenses for additional infrastructure and power consumption can be reduced to a minimum.

Prior work and state of the art. Among the previously proposed localization systems that use modulated LEDs, some employ a single photodiode in the receiver, while others rely on more elaborate receiver hardware. As an example of the second category, the receiver presented in [3] uses an image sensor and a lens, which allows it to determine the directions-of-arrival of light from multiple visible LEDs. The system proposed here, on the other hand, belongs to the first category, with the advantage of lower hardware complexity. Within this category, some receivers such as [4, 5] calculate their distance from the beacons based on the received signal strength, while others such as [6, 7] exploit the differences in the propagation delays from different beacons. All of the solutions cited above, however, rely on synchronization among the transmitters, which is costly to implement, while the proposed method does not require such synchronization. In [8], asynchronous transmission of the identifiers is enabled by applying the ALOHA protocol. In this setup, however, an identifier can only be transmitted successfully if no other beacon transmits at the same time, which may limit the time resolution. The method proposed here, on the other hand, allows all beacons to transmit their identifier simultaneously. Finally, all of the methods mentioned above rely on a deterministic signal model that assumes the presence of a line-of-sight to the beacons. The proposed method is based on a Bayesian model that takes into account the possibility of obstructions, thereby making it significantly more robust than methods ignoring this possibility.

Contributions. The main contributions of this paper are the following: First, due to the probabilistic signal model that takes into account disturbances, the proposed localization method is robust to obstructions of the line of sight from some of the beacons. Second, the proposed method does not require any synchronization among the beacons, thereby allowing them to be fully autonomous. Third, by approximating an optimal Bayesian estimator and taking into account the inherent sparsity in detecting the concurrently visible beacons, the proposed method achieves excellent performance while maintaining low complexity.

2. SYSTEM MODEL

A mobile device aims to determine its own position $\mathbf{z}$ (in two dimensions) within a building, using the visible light signals it receives from beacons in its vicinity. Each of the L beacons in the building is associated with an identifier, namely, a binary chip sequence $c^{(\ell)}[1], \dots, c^{(\ell)}[M]$ from $\{1, -1\}^M$, where $\ell \in \{1, \dots, L\}$ is the beacon index. The sequences are chosen such that both their autocorrelation for delays other than zero and their cross-correlation are low, which is ensured, e.g., in Gold sequences [9]. This usually implies that the sequence length M is larger than the number of beacons L. It is not necessary to ensure that each sequence appears at most

once within one building, because localization is never based on one single beacon. Each beacon modulates its light with the binary sequence using on-off-keying (where "on" and "off" correspond to 1 and -1, respectively). The sequence is repeated continuously. The chip rate $1/T$ is chosen high enough to ensure that the movement of the mobile device within the duration of the entire sequence MT can be neglected. This usually also ensures that the modulation is well above the frequency limit up to which the human eye would notice a flickering effect, which is a prerequisite. The sequences transmitted by different beacons are not synchronized with each other.

The mobile device knows the location of each beacon as well as the chip sequence associated with it. From its (unknown) position $\mathbf{z}$, only a subset $\mathcal{L}_z \subseteq \{1, \ldots, L\}$ of all beacons is potentially visible, whereas the other beacons are hidden behind walls. Normally, the number of potentially visible beacons $L_z = |\mathcal{L}_z|$ is much smaller than L. The mobile device knows the locations of walls, therefore it knows $\mathcal{L}_z$ as a function of $\mathbf{z}$. Using a single photodiode, the mobile device receives a superposition of the signals transmitted by several beacons and some additive noise. A simple amplifier circuit produces a signal that is proportional to the received intensity with its DC component (i.e., the time-average) removed. The signal is then sampled at rate N/T, with $N \in \mathbb{N}$. Let $r^{(\ell)}[k]$ denote the chip sequence $c^{(\ell)}[m]$ oversampled by a factor of N, i.e., $r^{(\ell)}[k] = c^{(\ell)}[m] \quad \text{for } k \in \big\{(m-1)N+1, ..., mN\big\}$. The sampled zero-average received signal $y[k]$ can then be expressed as

$$y[k] = \sum_{\ell \in \mathcal{L}_z} x_\ell \; r^{(\ell)}[k \oplus (-s_\ell)] + n[k]\,, \tag{1}$$

for $k = 1, \ldots, MN$. The contribution of beacon ℓ to $y[k]$ is weighted by a factor $x_\ell \in \mathbb{R}$ and circularly shifted by an offset $s_\ell \in \mathcal{S} = \{0, \ldots, MN-1\}$. The operator $\oplus$ denotes an addition combined with a modulo operation: $a \oplus b = \big((a+b-1) \bmod MN\big)+1$. For $\ell \notin \mathcal{L}_z$, we set $x_\ell = 0$ and $s_\ell = 0$. This allows us to define the vectors $\mathbf{x} = \big(x_1 \cdots x_L\big)^{\mathrm{T}}$ and $\mathbf{s} = \big(s_1 \cdots s_L\big)^{\mathrm{T}}$, where the superscript $^{\mathrm{T}}$ denotes transposition. Let $\mathbf{r}^{(\ell,s)}$ denote the length-MN vector that contains the sequence $r^{(\ell)}[k]$ shifted by s, i.e., $\mathbf{r}^{(\ell,s)} = \big(r^{(\ell)}[1 \oplus (-s)] \cdots r^{(\ell)}[MN \oplus (-s)]\big)^{\mathrm{T}}$. Using the matrix $\mathbf{R}^{(\mathbf{s})} = \big(\mathbf{r}^{(1,s_1)} \cdots \mathbf{r}^{(L,s_L)}\big)$, we can write (1) as follows:

$$\mathbf{y} = \mathbf{R}^{(\mathbf{s})}\,\mathbf{x} + \mathbf{n}\,, \tag{2}$$

where $\mathbf{y} = \big(y[1] \cdots y[MN]\big)^{\mathrm{T}}$ and $\mathbf{n} = \big(n[1] \cdots n[MN]\big)^{\mathrm{T}}$.

The weight factor x_ℓ for $\ell \in \mathcal{L}_z$ depends on the distance between the position $\mathbf{z}$ of the mobile device and the position $\mathbf{p}_\ell$ of beacon ℓ. We model x_ℓ as a product of two factors:

$$x_\ell \;=\; h_\ell\, A(d_{\mathbf{z},\ell}) \quad \text{with} \quad d_{\mathbf{z},\ell} = \|\mathbf{p}_\ell - \mathbf{z}\|\,.$$

Here, h_ℓ is a stochastic shadowing factor that represents the probabilities of some obstructions blocking the line of sight or attenuating the signal strength. The second factor $A(d_{\mathbf{z},\ell})$ represents the received signal strength in the absence of obstructions. Under some simplifying assumptions, the expression presented in [10] for line-of-sight links and generalized Lambertian transmitters leads to the following relation:

$$A(d_{\mathbf{z},\ell}) = \frac{A_{0,\ell}\, d_{v,\ell}^2}{\big(d_{v,\ell}^2 + d_{\mathbf{z},\ell}^2\big)^2}\,, \tag{3}$$

where $d_{v,\ell}$ is the vertical distance between the transmitter and the receiver and $A_{0,\ell}$ (which depends on the brightness of beacon lamp ℓ, among other things) is constant with respect to $d_{\mathbf{z},\ell}$. The mobile device knows $d_{v,\ell}$ and $A_{0,\ell}$ for all ℓ. For practical application of the proposed method, the the function $A(d_{\mathbf{z},\ell})$ should be adapted according to the characteristics of the transmitting and receiving diodes as well as the local environment. The functionality of the proposed method is not constrained to a particular choice of $A(d_{\mathbf{z},\ell})$.

The time-shifts s_ℓ contain no useful information, since the beacons and the mobile device are not synchronized. While the time-shifts s_ℓ in our model are discrete, the true propagation delays between beacons and the mobile device are not. This means that there is a small but nonzero probability that some sampling instants of the receiver fall into the short period while the received signal from beacon ℓ switches between -1 and 1. However, at most $1/N$ of the samples may be affected, which justifies the use of larger N, at the cost of higher complexity. The residual effect can be modeled as an effect of the stochastic shadowing factor h_ℓ.

3. STOCHASTIC MODEL

For the sake of simplicity, the additive noise $\mathbf{n}$ is assumed to be zero-mean, white, and Gaussian with variance σ_n^2. The likelihood function $p(\mathbf{y}|\mathbf{x}, \mathbf{s})$ is therefore given as

$$p(\mathbf{y}|\mathbf{x}, \mathbf{s}) = \mathcal{N}\big(\mathbf{R}^{(\mathbf{s})}\,\mathbf{x}, \sigma_n^2 \mathbf{I}\big)\,, \tag{4}$$

where $\mathbf{I}$ denotes the identity matrix. In practice, the role of additive noise in visible light channels is often less significant than other effects such as interference between beacons, loss of the line of sight, or random fluctuations of the received signal strength.

In our model, the parameters $\mathbf{z}$, h_ℓ, and s_ℓ (both for all $\ell \in \mathcal{L}_z$) are random. The location $\mathbf{z}$ is assigned a uniform prior:

$$p(\mathbf{z}) = 1/D\,, \tag{5}$$

where D is the area within which we intend to localize the mobile device. The prior of the shadowing factor h_ℓ is defined as the weighted sum of two Gaussian distributions. One of them represents the case that the line of sight is not obstructed and the received signal strength is close to its theoretical value $A(d_{\mathbf{z},\ell})$; this Gaussian distribution is centered around 1 and weighted with $\alpha \in [0, 1]$. The other one represents the case that the line of sight is obstructed or the signal is attenuated; it has a larger variance and is centered around 0 to describe weak indirect illumination from the beacon. This yields

$$p(h_\ell) = \alpha\,\mathcal{N}\big(1, \sigma_1^2\big) + (1-\alpha)\,\mathcal{N}\big(0, \sigma_2^2\big)\,.$$

It follows that the prior of x_ℓ given the position $\mathbf{z}$ is

$$p(x_\ell|\mathbf{z}) = \frac{\alpha}{\sqrt{2\pi}\sigma_1 A(d_{\mathbf{z},\ell})}\,\exp\left(-\frac{\big(x_\ell - A(d_{\mathbf{z},\ell})\big)^2}{2\big(\sigma_1 A(d_{\mathbf{z},\ell})\big)^2}\right) + \frac{1-\alpha}{\sqrt{2\pi}\sigma_2 A(d_{\mathbf{z},\ell})}\,\exp\left(-\frac{x_\ell^2}{2\big(\sigma_2 A(d_{\mathbf{z},\ell})\big)^2}\right), \tag{6}$$

for $\ell \in \mathcal{L}_z$. The shift s_ℓ for $\ell \in \mathcal{L}_z$ is assigned a uniform prior:

$$p(s_\ell|\mathbf{z}) = 1/(MN)\,. \tag{7}$$

For $\ell \notin \mathcal{L}_z$, x_ℓ and s_ℓ are zero, i.e., $p(x_\ell|\mathbf{z}) = \delta(x_\ell)$ and $p(s_\ell|\mathbf{z}) = \delta(s_\ell)$. For given $\mathbf{z}$, the parameters $\mathbf{x}$ and $\mathbf{s}$ are *a priori* statistically independent from each other and from the noise $\mathbf{n}$, and their elements are also statistically independent from each other.

4. LOCALIZATION

4.1. MAP estimation of z

Our goal is to estimate the position $\mathbf{z}$, given the observed signal $\mathbf{y}$. While the true position $\mathbf{z}$ is not discrete by nature, we discretize our estimate by choosing it from a set $\mathcal{Z}$ that contains Z hypotheses. The *maximum a posteriori* (MAP) detector is theoretically optimal

in the sense that it minimizes the probability of an incorrect detection if the true $\mathbf{z}$ is also from $\mathcal{Z}$ (which is, however, not the case here). The MAP detector is found by maximizing the posterior distribution $p(\mathbf{z}|\mathbf{y})$, which can be obtained by marginalizing the joint posterior $p(\mathbf{z},\mathbf{x},\mathbf{s}|\mathbf{y})$:

$$\begin{aligned}\hat{\mathbf{z}}_{\text{MAP}} &= \underset{\mathbf{z}}{\operatorname{argmax}}\; p(\mathbf{z}|\mathbf{y}) = \underset{\mathbf{z}}{\operatorname{argmax}} \sum_{\mathbf{s}} \int p(\mathbf{z},\mathbf{x},\mathbf{s}|\mathbf{y})\,\mathrm{d}\mathbf{x} \\ &= \underset{\mathbf{z}}{\operatorname{argmax}} \sum_{\mathbf{s}} \int \frac{p(\mathbf{y}|\mathbf{x},\mathbf{s})\,p(\mathbf{x}|\mathbf{z})\,p(\mathbf{z})\,p(\mathbf{s}|\mathbf{z})}{p(\mathbf{y})}\,\mathrm{d}\mathbf{x}\,,\end{aligned}$$

where Bayes' rule has been used. By dropping factors that are constant with respect to $\mathbf{z}$, $\mathbf{x}$, and $\mathbf{s}$ (cf. (5) and (7)), this simplifies to

$$\hat{\mathbf{z}}_{\text{MAP}} = \underset{\mathbf{z}}{\operatorname{argmax}} \sum_{\mathbf{s}} \int p(\mathbf{y}|\mathbf{x},\mathbf{s})\,p(\mathbf{x}|\mathbf{z})\,\mathrm{d}\mathbf{x}\,. \tag{8}$$

In practice, the summation over all $\mathbf{s} \in \mathcal{S}^L = \{0,\dots,MN-1\}^L$ simplifies to a sum over $\mathcal{S}^{L_z}$, since only L_z elements of $\mathbf{s}$ can be nonzero. Nevertheless, the complexity of (8) is still prohibitive. We resort to the following modification: instead of marginalizing $p(\mathbf{z},\mathbf{x},\mathbf{s}|\mathbf{y})$ with respect to $\mathbf{s}$, we insert an estimate $\hat{\mathbf{s}}$ and thus obtain (after simplifications analogous to those above)

$$\hat{\mathbf{z}} = \underset{\mathbf{z}}{\operatorname{argmax}} \int p(\mathbf{y}|\mathbf{x},\hat{\mathbf{s}})\,p(\mathbf{x}|\mathbf{z})\,\mathrm{d}\mathbf{x}\,. \tag{9}$$

The estimate $\hat{\mathbf{s}}$ is obtained from the following joint estimator of $\mathbf{x}$ and $\mathbf{s}$:

$$(\hat{\mathbf{x}}_{\text{CS}},\hat{\mathbf{s}}_{\text{CS}}) = \underset{(\mathbf{x},\mathbf{s})}{\operatorname{argmin}} \left\|\mathbf{y}-\mathbf{R}^{(\mathbf{s})}\,\mathbf{x}\right\|^2 + \lambda\,\|\mathbf{x}\|_1\,, \tag{10}$$

where $\|\cdot\|_1$ stands for the ℓ_1-norm and $\lambda > 0$. This joint estimator is discussed in Section 4.2. The estimator of $\mathbf{z}$ in (9) is of tractable complexity and has shown good performance in our simulations. However, by using the second output of the time-shift estimator (i.e., $\hat{\mathbf{x}}_{\text{CS}}$), we can achieve another dramatic reduction of the complexity of (9). We start by finding the relation of $\hat{\mathbf{x}}_{\text{CS}}$ and the true $\mathbf{x}$. Let $\hat{\mathcal{L}}$ denote a set containing the indices of the nonzero elements of $\hat{\mathbf{x}}_{\text{CS}}$, and let $\hat{\mathbf{x}}_{\hat{\mathcal{L}}}$, $\mathbf{x}_{\hat{\mathcal{L}}}$, and $\mathbf{R}^{(\mathbf{s})}_{\hat{\mathcal{L}}}$ comprise the corresponding elements of $\hat{\mathbf{x}}_{\text{CS}}$ and $\mathbf{x}$ and the corresponding columns of $\mathbf{R}^{(\mathbf{s})}$, respectively. Furthermore, we use $\mathbf{A}^\sharp$ to denote the left pseudo-inverse of a matrix $\mathbf{A}$, i.e., $\mathbf{A}^\sharp = \left(\mathbf{A}^{\mathrm{T}}\mathbf{A}\right)^{-1}\mathbf{A}^{\mathrm{T}}$. It can be shown that because of (10), $\hat{\mathbf{x}}_{\hat{\mathcal{L}}}$ can be expressed as

$$\begin{aligned}\hat{\mathbf{x}}_{\hat{\mathcal{L}}} &= \left(\mathbf{R}^{(\hat{\mathbf{s}}_{\text{CS}})}_{\hat{\mathcal{L}}}\right)^\sharp \mathbf{y} = \left(\mathbf{R}^{(\hat{\mathbf{s}}_{\text{CS}})}_{\hat{\mathcal{L}}}\right)^\sharp \left(\mathbf{R}^{(\mathbf{s})}\,\mathbf{x}+\mathbf{n}\right) \\ &= \mathbf{x}_{\hat{\mathcal{L}}} + \underbrace{\left(\mathbf{R}^{(\hat{\mathbf{s}}_{\text{CS}})}_{\hat{\mathcal{L}}}\right)^\sharp \left(\mathbf{R}^{(\mathbf{s})}\,\mathbf{x} - \mathbf{R}^{(\hat{\mathbf{s}}_{\text{CS}})}_{\hat{\mathcal{L}}}\,\mathbf{x}_{\hat{\mathcal{L}}}\right) + \left(\mathbf{R}^{(\hat{\mathbf{s}}_{\text{CS}})}_{\hat{\mathcal{L}}}\right)^\sharp \mathbf{n}}_{\mathbf{v}}\,,\end{aligned}$$

where the vector $\mathbf{v}$ comprises the filtered noise as well as distortions due to erroneous $\hat{\mathbf{s}}_{\text{CS}}$ and $\hat{\mathcal{L}}$. Simulation results have shown that the statistical dependence of $\mathbf{v}$ on $\mathbf{x}$ and $\mathbf{s}$ can be neglected in our estimation method without compromising its performance. This seems plausible, since the columns of $\mathbf{R}^{(\hat{\mathbf{s}}_{\text{CS}})}$ are only weakly correlated, on one hand, and significant errors in $\hat{\mathbf{s}}_{\text{CS}}$ are very rare, on the other hand. We therefore model $\mathbf{v}$ as white and Gaussian with variance

$$\sigma_v^2 = \tfrac{1}{|\hat{\mathcal{L}}|}\,\sigma_n^2\,\mathrm{tr}\big(\mathbf{P}^{-1}\big) \quad \text{with} \quad \mathbf{P} = \left(\mathbf{R}^{(\hat{\mathbf{s}}_{\text{CS}})}_{\hat{\mathcal{L}}}\right)^{\mathrm{T}} \mathbf{R}^{(\hat{\mathbf{s}}_{\text{CS}})}_{\hat{\mathcal{L}}}\,, \tag{11}$$

where $|\hat{\mathcal{L}}|$ denotes the cardinality of $\hat{\mathcal{L}}$. The same relation is adopted for the remaining elements of $\hat{\mathbf{x}}_{\text{CS}}$ and $\mathbf{x}$, i.e.,

$$p(\hat{\mathbf{x}}_{\text{CS}}|\mathbf{x}) = \mathcal{N}\big(\mathbf{x},\sigma_v^2\mathbf{I}\big)\,. \tag{12}$$

Using this distribution as a likelihood function, we can reformulate our estimator of $\mathbf{z}$ based on $\hat{\mathbf{x}}_{\text{CS}}$ as observed data, rather than on $\mathbf{y}$. The key advantage of the likelihood function (12) compared to (4), besides its lower dimensionality, is that it can be factorized as follows: $p(\hat{\mathbf{x}}_{\text{CS}}|\mathbf{x}) = \prod_{\ell=1}^{L} p(\hat{x}_\ell|x_\ell)$, where each factor depends only on one element of $\mathbf{x}$. By replacing $p(\mathbf{y}|\mathbf{x},\hat{\mathbf{s}})$ with $p(\hat{\mathbf{x}}_{\text{CS}}|\mathbf{x})$ in (9), we obtain

$$\begin{aligned}\hat{\mathbf{z}} &= \underset{\mathbf{z}}{\operatorname{argmax}} \int p(\hat{\mathbf{x}}_{\text{CS}}|\mathbf{x})\,p(\mathbf{x}|\mathbf{z})\,\mathrm{d}\mathbf{x} \\ &= \underset{\mathbf{z}}{\operatorname{argmax}} \int \prod_{\ell=1}^{L} p(\hat{x}_\ell|x_\ell)p(x_\ell|\mathbf{z})\,\mathrm{d}\mathbf{x} \\ &= \underset{\mathbf{z}}{\operatorname{argmax}} \prod_{\ell=1}^{L} \int e^{-(\hat{x}_\ell-x_\ell)^2/(2\sigma_v^2)}\,p(x_\ell|\mathbf{z})\,\mathrm{d}x_\ell\,.\end{aligned}$$

In the last step, we used (12) and dropped some constant factors. It can be shown that inserting (6) finally leads to

$$\hat{\mathbf{z}} = \underset{\mathbf{z}}{\operatorname{argmax}} \prod_{\ell\in\mathcal{L}_z} \left(\eta_\ell\, e^{\varphi_\ell^2(\hat{x}_\ell+\gamma_\ell)^2} + \xi_\ell\, e^{\psi_\ell^2\hat{x}_\ell^2}\right), \tag{13}$$

with

$$\eta_\ell = \frac{\alpha\,\varphi_\ell\,\sqrt{2}\sigma_v^2}{\sigma_1\,A(d_{\mathbf{z},\ell})\,e^{1/(2\sigma_1^2)}}\,, \qquad \gamma_\ell = \frac{\sigma_v^2}{\sigma_1^2 A(d_{\mathbf{z},\ell})}\,,$$

$$\varphi_\ell = \left(\frac{2\sigma_v^4}{\sigma_1^2\big(A(d_{\mathbf{z},\ell})\big)^2} + 2\sigma_v^2\right)^{-\frac{1}{2}}, \quad \xi_\ell = \frac{(1-\alpha)\,\psi_\ell\,\sqrt{2}\sigma_v^2}{\sigma_2\,A(d_{\mathbf{z},\ell})}\,,$$

and ψ_ℓ defined analogously to φ_ℓ by replacing σ_1^2 with σ_2^2. After observing $\mathbf{y}$, the proposed localization method first estimates $\mathbf{x}$ and $\mathbf{s}$ as described in Section 4.2 and then estimates $\mathbf{z}$ according to (13). To further reduce complexity (without changing $\hat{\mathbf{z}}$), the set of all hypotheses $\mathcal{Z}$ can be replaced by the set of all positions from which at least one of the beacons in $\hat{\mathcal{L}}$ can be seen.

4.2. Joint estimation of $\mathbf{x}$ and $\mathbf{s}$

Estimating both $\mathbf{x}$ and $\mathbf{s}$ from $\mathbf{y}$ (according to the system model in (2)) is an ill-posed problem. To see this more clearly, we extend the definition of x_ℓ to a 2-dimensional domain comprising all beacons and all possible shifts, i.e., $\chi_{\ell,s} = x_\ell$ if $s = s_\ell$ and $\chi_{\ell,s} = 0$ otherwise. Using this definition, we can rewrite (2) as follows:

$$\mathbf{y} \;=\; \mathbf{R}\,\boldsymbol{\chi} + \mathbf{n}\,, \tag{14}$$

with the length-LMN vector $\boldsymbol{\chi} = \big(\chi_{1,1}\;\cdots\;\chi_{1,MN}\;\chi_{2,1}\;\cdots\;\chi_{L,MN}\big)^{\mathrm{T}}$ and the $MN \times LMN$ matrix $\mathbf{R} = \big(\mathbf{r}^{(1,1)}\;\cdots\;\mathbf{r}^{(1,MN)}\;\mathbf{r}^{(2,1)}\;\cdots\;\mathbf{r}^{(L,MN)}\big)$. Now $\mathbf{R}$ does not depend on $\mathbf{s}$, while $\boldsymbol{\chi}$ contains all information about both $\mathbf{x}$ and $\mathbf{s}$. The formulation in (14) clearly shows that the problem of estimating $\boldsymbol{\chi}$ from $\mathbf{y}$ is ill-posed, since the number of measurements MN is much smaller than the number of signal dimensions LMN. In this situation, a classical approach would be to estimate each (x_ℓ, s_ℓ) separately, by maximizing the output of a matched filter with the corresponding sequence $r^{(\ell)}[k]$:

$$\hat{x}_{\ell,\text{MF}} = \max_{s\in\mathcal{S}} \big(\mathbf{r}^{(\ell,s)}\big)^{\mathrm{T}}\,\mathbf{y}\,, \quad \hat{s}_{\ell,\text{MF}} = \underset{s\in\mathcal{S}}{\operatorname{argmax}} \big(\mathbf{r}^{(\ell,s)}\big)^{\mathrm{T}}\,\mathbf{y}\,. \tag{15}$$

These straight-forward estimators are easy to implement, but they ignore the correlation among different sequences $\mathbf{r}^{(\ell,s)}$, which limits their performance. Since the signal strength declines rapidly with increasing distance (cf. (3)), even weak correlation causes strong interference from the closer beacons in the signals from more distant beacons.

However, the particular properties of $\boldsymbol{\chi}$ allow for joint estimation of $(\mathbf{x}, \mathbf{s})$, which overcomes this problem. Note that $\boldsymbol{\chi}$ is highly sparse; only L_z of its LMN elements are nonzero. Besides sparsity, $\boldsymbol{\chi}$ must also satisfy an additional structural constraint: for each ℓ, only one of the MN elements $\chi_{\ell,s}$ may be nonzero. This can be formulated as a block constraint: within $\boldsymbol{\chi}$, each block of length MN may contain at most one nonzero element. Let $\mathcal{C} \subseteq \mathbb{R}^{LMN}$ denote the set of all vectors $\boldsymbol{\chi}$ that satisfy this constraint. Then (10) can be written as

$$\hat{\boldsymbol{\chi}}_{\text{CS}} = \underset{\boldsymbol{\chi} \in \mathcal{C}}{\operatorname{argmin}} \, \|\mathbf{y} - \mathbf{R}\,\boldsymbol{\chi}\|^2 + \lambda \, \|\boldsymbol{\chi}\|_1 \, .$$

This can be seen as a problem of sparse-signal recovery. Corresponding methods have been studied extensively in the field of compressive sensing (CS) [11], including localization and radar applications [12]. Due to its sparsity, the vector $\boldsymbol{\chi}$ can be recovered stably if the measurement matrix $\mathbf{R}$ satisfies certain properties [11]. A popular algorithm for solving such problems is orthogonal matching pursuit (OMP) [13]. Algorithm 1 shows the steps of the proposed OMP algorithm for estimating $\mathbf{x}$ and $\mathbf{s}$. Here, $\hat{\mathbf{s}} = \big(\hat{s}_1 \cdots \hat{s}_L\big)^{\text{T}}$, $\mathbf{R}_{\mathcal{W}^{(i)}}$ denotes the $MN \times i$ matrix whose columns are $\mathbf{r}^{(\ell,s)}$ for all $(\ell, s) \in \mathcal{W}^{(i)}$, and $\varrho > 0$. In our simulations, $\varrho = 0.85\, MN\sigma_n^2$ showed to be a good choice.

Algorithm 1 OMP algorithm for estimating $\mathbf{x}$ and $\mathbf{s}$

1: $\tilde{\mathbf{y}}^{(0)} \leftarrow \mathbf{y}$, $\mathcal{L}^{(0)} \leftarrow \{1, \ldots, L\}$, $\mathcal{W}^{(0)} \leftarrow \emptyset$, $\hat{\mathbf{x}}_{\text{CS}} \leftarrow \mathbf{0}$, $\hat{\mathbf{s}} \leftarrow \mathbf{0}$
2: Iterate for $i = 1, 2, \ldots$ while $i \leq L$ and $\|\tilde{\mathbf{y}}^{(i-1)}\|^2 > \varrho$
3: $\quad \big(\ell^{(i)}, s^{(i)}\big) \leftarrow \operatorname{argmax}_{(\ell,s) \in \mathcal{L}^{(i-1)} \times \mathcal{S}} \, \big|\big(\mathbf{r}^{(\ell,s)}\big)^{\text{T}} \tilde{\mathbf{y}}^{(i-1)}\big|$
4: $\quad \mathcal{L}^{(i)} \leftarrow \mathcal{L}^{(i-1)} \backslash \ell^{(i)}$
5: $\quad \hat{s}_{\ell^{(i)}} \leftarrow s^{(i)}$
6: $\quad \mathcal{W}^{(i)} \leftarrow \mathcal{W}^{(i-1)} \cup \big(\ell^{(i)}, s^{(i)}\big)$
7: $\quad \mathbf{x}^{(i)} \leftarrow \operatorname{argmin}_{\mathbf{x} \in \mathbb{R}^i} \|\mathbf{y} - \mathbf{R}_{\mathcal{W}^{(i)}}\, \mathbf{x}\|^2$
8: $\quad \tilde{\mathbf{y}}^{(i)} \leftarrow \mathbf{y} - \mathbf{R}_{\mathcal{W}^{(i)}}\, \mathbf{x}^{(i)}$
9: $\hat{\mathbf{x}}_{\hat{\mathcal{L}}} \leftarrow \mathbf{x}^{(i)}$, $\hat{\mathbf{s}}_{\text{CS}} \leftarrow \hat{\mathbf{s}}$

In the unlikely case that two beacons ℓ_1 and ℓ_2 use the same chip sequence, we can still assume that the two corresponding time-shifts are different. Then, the only effect of the identical sequences is that the two elements of $\hat{\mathbf{x}}_{\text{CS}}$ corresponding ℓ_1 and ℓ_2 may be switched. We can cope with this by calculating $\hat{\mathbf{z}}$ for each version of $\hat{\mathbf{x}}_{\text{CS}}$ and choosing the one for which the decision metric of (13) is larger.

5. NUMERICAL RESULTS

To evaluate the performance and complexity of the proposed method, we performed 10000 simulations using the floorplan shown in Fig. 1. In the following, all lengths are normalized with respect to a unit length of 1 meter. We randomly generated the parameters according to the priors given in Section 3, with $L = 137$, $M = 256$, $N = 4$, $A_{0,\ell} = 2.25$, $d_{v,\ell} = 1.5$ for all ℓ, $\alpha = 0.7$, $\sigma_1^2 = 2 \cdot 10^{-2}$, and $\sigma_2^2 = 4.6 \cdot 10^{-2}$. The chip sequences $c^{(\ell)}[m]$ were randomly generated in each simulation. To obtain a constant signal-to-noise ratio of 10dB, we set $\sigma_n^2 = \|\mathbf{R}^{(\mathbf{s})}\, \mathbf{x}\|^2/(10MN)$ in each simulation. The estimator calculates σ_v^2 according to (11). The $Z = 14400$ hypotheses for $\hat{\mathbf{z}}$ are arranged in a rectangular grid of dimensions 120×120 with minimum distance 0.25 (i.e., 25cm). The true positions of the mobile device $\mathbf{z}$ are not discretized to this grid but drawn from a (spatially continuous) uniform distribution over the entire area of the floorplan, as stated in (5).

We compare the results of our method to those of two reference methods. The methods listed in Section 1 are not designed for this kind of setup and can thus not easily be compared. Instead, we resort to a classical matched-filter approach as a reference method (labeled "RM1"). In the first step, each $\hat{x}_\ell$ is obtained as the maximum output of a time-domain matched filter according to (15). In the second step, $\hat{\mathbf{z}}$ is obtained as the maximizer of a "beacon-domain" matched filter: $\hat{\mathbf{z}}_{\text{MF}} = \operatorname{argmax}_{\mathbf{z}} \sum_{\ell=1}^{L} \hat{x}_\ell \, A(d_{\mathbf{z},\ell})$. The second reference method (labeled "RM2") first estimates $\hat{\mathbf{x}}_{\text{CS}}$ according to (10) like the proposed method and then uses it to calculate $\hat{\mathbf{z}}_{\text{MF}}$ like RM1. Fig. 2 shows the empirical cumulative distribution function (cdf) of the localization error $\epsilon = \|\hat{\mathbf{z}} - \mathbf{z}\|$ obtained with the three methods. The proposed method clearly outperforms the reference methods in terms of the localization error. The root mean squared error of $\hat{\mathbf{z}}$ obtained with the proposed method is 0.81, while that of RM1 and RM2 is 2.85 and 1.81, respectively (recall that 1 corresponds to a length of 1m). The average computation time of the proposed method for an unoptimized MATLAB R2012a 64-bit implementation on a 2.93-GHz Intel Core i7-870 processor (which may be reduced significantly when implemented in practice) was 0.33s, which seems fair enough for practical applications.

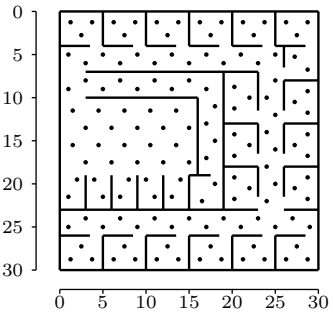

Fig. 1. Floor plan of dimensions 30×30 (with unit length 1m). The dots represent the locations of beacons.

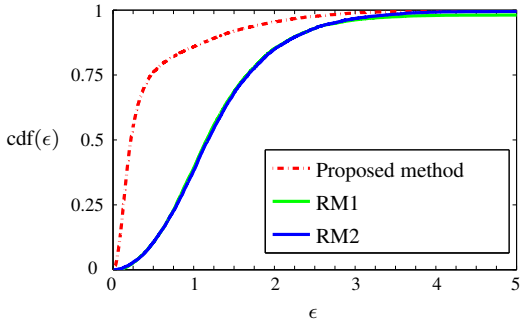

Fig. 2. Empirical cumulative distribution function of $\epsilon = \|\hat{\mathbf{z}} - \mathbf{z}\|$ obtained in 10000 simulations. For $\epsilon < 3.5$, the lines of RM1 and RM2 almost coincide.

6. CONCLUSION

We proposed a Bayesian method for indoor self-localization of mobile devices using visible light from modulated LED light sources. In contrast to existing systems, the proposed method does not require synchronization of the transmitters. This allows the light sources to be fully autonomous, ensuring a minimum of infrastructure requirements and costs. By using a suitable probabilistic model and taking into account the inherent sparsity in detecting the concurrently active light sources, the proposed method achieves high robustness with respect to obstructions of the lines of sight, excellent performance, and low complexity.

7. REFERENCES

[1] W. Zhang and M. Kavehrad, "Comparison of VLC-based indoor positioning techniques," in *Proc. SPIE 8645, Broadband Access Commun. Technol. VII*, 2013, pp. 86450M/1–6.

[2] H. Liu, H. Darabi, P. Banerjee, and J. Liu, "Survey of wireless indoor positioning techniques and systems," *IEEE Trans. Systems, Man, Cybernetics, Part C: Applic. and Reviews*, vol. 37, no. 6, pp. 1067–1080, 2007.

[3] M. Yoshino, S. Haruyama, and M. Nakagawa, "High-accuracy positioning system using visible LED lights and image sensor," in *Proc. IEEE Radio Wireless Symp.*, Orlando, FL, USA, Jan. 2008, pp. 439–442.

[4] S. Hann, J.-H. Kim, S.-Y. Jung, and C.-S. Park, "White LED ceiling lights positioning systems for optical wireless indoor applications," in *Proc. 36th Europ. Conf. and Exhib. Optical Commun. (ECOC)*, Torino, Italy, Sept. 2010, pp. 1–3.

[5] Z. Zhou, M. Kavehrad, and P. Deng, "Indoor positioning algorithm using light-emitting diode visible light communications," *Optical Engineering*, vol. 51, no. 8, pp. 085009/1–6, 2012.

[6] K. Panta and J. Armstrong, "Indoor localisation using white LEDs," *Electronics letters*, vol. 48, no. 4, pp. 228–230, 2012.

[7] S.-Y. Jung, S. Hann, and C.-S. Park, "TDOA-based optical wireless indoor localization using LED ceiling lamps," *IEEE Trans. Consumer Electronics*, vol. 57, no. 4, pp. 1592–1597, 2011.

[8] W. Zhang and M. Kavehrad, "A 2-D indoor localization system based on visible light LED," in *Proc. IEEE Photonics Society Summer Topical Meeting Series*, Seattle, WA, USA, July 2012, pp. 80–81.

[9] R. Gold, "Optimal binary sequences for spread spectrum multiplexing (Corresp.)," *IEEE Trans. Information Theory*, vol. 13, no. 4, pp. 619–621, 1967.

[10] J.M. Kahn and J.R. Barry, "Wireless infrared communications," *Proc. IEEE*, vol. 85, no. 2, pp. 265–298, 1997.

[11] E.J. Candés and M.B. Wakin, "An introduction to compressive sampling," *IEEE Sign. Process. Mag.*, vol. 25, no. 2, pp. 21–30, 2008.

[12] P. Maechler, N. Felber, and H. Kaeslin, "Compressive sensing for WiFi-based passive radar," in *Proc. Eusipco*, Bucharest, Romania, Aug. 2012, pp. 1444–1448.

[13] Y.C. Pati, R. Rezaiifar, and P.S. Krishnaprasad, "Orthogonal matching pursuit: recursive function approximation with applications to wavelet decomposition," in *Proc. Asilomar Conf. Signals, Systems, Computers*, Pacific Grove, CA, USA, Nov. 1993, pp. 40–44.