# 5 Electronic Locomotion Aids for the Blind: Towards More Assistive Systems

Christophe Jacquet[1,2], Yacine Bellik[2], and Yolaine Bourda[1]

[1] Supélec, Plateau de Moulon, 91192 Gif-sur-Yvette Cedex, France
Christophe.Jacquet@supelec.fr, Yolaine.Bourda@supelec.fr
[2] LIMSI-CNRS, BP 133, 91403 Orsay Cedex, France
Yacine.Bellik@limsi.fr

**Abstract.** This chapter first presents a review of existing locomotion assistance devices for the blind. These devices are merely proximeters, that measure the distance to the closest obstacles, and convey this information to their users. We introduce the measurement methods (infrared sensors, ultrasonic sensors, laser telemeters) and the user interfaces (sounds and tactile vibrations). Then, we analyse the shortcomings of these systems, and thus explain what additional features new devices could offer. To study the feasibility of such systems, we tackle the different issues raised in the process: localizing users, modeling their environment and adding semantic annotations. Finally, we explain how such devices could fit into a view of ambient intelligence, and how the problems raised extend beyond the field of assistance to blind people.

## 1 Introduction

According to the World Health Organization, there are 45 million blind people in the world, which amounts to an estimated 1‰ to 2‰ of the population in industrialized countries. This figure cannot be neglected, and the problems encountered by blind people in their everyday life need to be addressed.

In particular, these people are faced with huge difficulties moving in cities, where streets, public transportation systems and shopping malls represent hostile ever-changing environments. As a result, blind people are in danger while moving on their own, and their autonomy is limited.

Indeed, if blind people can generally remember their way to some places, they cannot know in advance what obstacles they will stumble upon. In consequence, the fear of the unknown often leads them to restrict their universe to a small set of known places. They do not dare going anywhere else, thus experiencing limited travel freedom.

Leader dogs can help blind people avoid obstacles and find their way in unknown environments, but they are very expensive: the cost per blind person – leader dog pair ranges from $15,000 to $30,000. Despite the financial support of some organizations such as the Lions Club, only very few blind people can actually have a leader dog, while many of them would like to.

However, leader dogs represent a very valuable means of dealing with unknown places. Indeed, they see obstacles at a distance, and can therefore anticipate the necessary avoidance maneuvers. In contrast, blind people using canes can feel silent obstacles at a cane's length distance only, so their anticipation capability is very poor [1]. The dog can help them improve their anticipation performances, which leads directly to more fluent trajectories, more self-confidence and easier travel [2].

This is the reason why numerous electronic locomotion assistance systems have been developed. The vast majority of existing systems are obstacle detectors that can warn users of the presence of obstacles *in advance*, so as to enable them to *anticipate* the presence of obstacles and adapt their behavior accordingly. Electronic devices are significantly cheaper than leader dogs, although offering more limited, yet valuable information.

It is important to notice that the white cane is not only a useful obstacle detector for blind people, but also a means by which blind people are recognized by sighted people. In consequence, all locomotion assistance devices for the blind must be designed to be secondary aids, used in complement to the long cane, and not instead of it, because they are not inherently social indicators of blindness.

In this chapter, we first present an overview of existing obstacle detectors, while distinguishing two different facets in these systems: information capture and information presentation. After, we discuss the shortcomings of existing systems and we draw a list of requirements for more advanced systems. Then, we give details about the implementation issues of this new class of systems: user localization and environment modeling, from the structural and semantic points of view. Finally, we show how these issues take place within the broader field if ambient intelligence.

## 2 Existing techniques

Navigation assistance devices need to deal with two different issues. First, they need to capture contextual information, in general distance information. Indeed, because the basic purpose of these devices is to warn users of approaching obstacles, their basic requirement is to measure distances to obstacles.

Second, they need to present users with this information. The presentation method must be adapted to blind users, and must be suitable for continuous use.

Several electronic travel aids have been proposed that address both of these problems. We have studied the characteristics of a few of them that we think are representative of the available techniques. Below, we present the two problems and the solutions proposed by these devices.

### 2.1 Information Capture

All current devices only take distance information into account. Thus, they use various kinds of telemeters and proximeters.

**Infrared Sensors** Infrared sensors are characterized by their quite wide angle magnitude (20°) which allows the detection of obstacles in the general heading direction of the user. However, their range is limited to a few meters, so they cannot warn users very long before encountering an obstacle.

Tom Pouce (fig. 1), developed by the LAC[3], is an infrared proximeter that can detect obstacles within a 0.5, 1.5 or 3 meters range. Users can select the desired range by means of a three-position switch. Thus they can detect either very close or more remote objects, depending on their current task, for instance following a corridor or trying to find a door in a wall.

**Fig. 1.** Photo of the Tom Pouce.

**Ultrasonic Sensors** Ultrasonic sensors have roughly the same characteristics as the infrared sensors, in terms of range and angle magnitude.

The Miniguide[4], the Polaron[5] and the UltraCane[6] (fig. 2) are examples of mobility aids that use ultrasonic sensors to detect obstacles.

Both devices can be configured to detect obstacles in various distance ranges. The range can be set to 0.5, 1, 2 or 4 meters for the Miniguide, and to 1.2, 2.4 or 4.8 meters for the Polaron.

The Miniguide has got two sensors that should be one above the other while in operation. Thus, users must pay attention to hold their devices vertically.

The Polaron avoids this shortcoming thanks to a neck strap that allows users to wear it around the neck at chest height, thus ensuring a constant ideal position.

---

[3] Laboratoire Aimé Cotton, Orsay, France.

[4] The Miniguide is a product of the Australian company GDP Research, see `http://www.gdp-research.com.au/`.

[5] The Polaron is a product of the company Nurion-Raycal, see `http://www.nurion.net/`.

[6] The UltraCane, previously known as "Batcane", is a product of the company Sound Foresight, see `http://www.soundforesight.co.uk/`.

**Fig. 2.** Photos of the Miniguide (top left), the Ultracane (bottom left) and the Polaron (right).

**Laser Telemeters** One good example of laser telemeter is the Teletact 1, developed by the LAC.

The distance to the obstacle encountered by the laser beam is measured with about 1% accuracy in the 10 cm – 10 m range. Fig. 3 gives the basic principle of the telemeter [3], and fig. 4 is a photograph of the commercial device. Blind users perceive the direction they point at thanks to the internal awareness of the position of their limbs (proprioception).

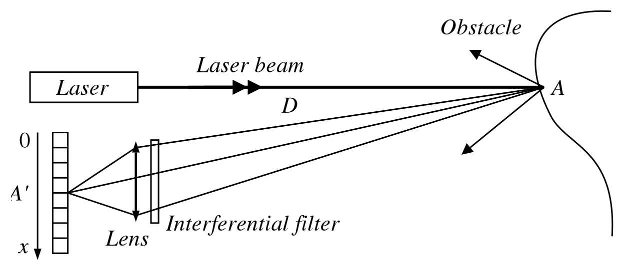


**Fig. 3.** Basic principle of the Teletact 1's telemeter.

In the device, a laser diode emits a red (670 nm) laser beam that can range up to 30 m (but measurements are precise enough up to 10 m only). For ocular security, the laser power is 1 mW (class II laser product). Thus, there is no risk for passers-by, even if the blind user points the beam at their eyes by mistake.

The beam meets the obstacle at a distance $D$ and creates a laser spot $A$. The image of the laser spot $A$ through the lens on the CCD line is $A'$. The position of $A'$ on the line gives the distance $D$. Actually, $x_{A'}$ is inversely proportional to $D$, which leads to very simple calculations.

The basic principle is straightforward; the main difficulty is to achieve immunity to diurnal light (up to 200 klux on sunny summer days) that requires the use of some additional optical and electronic systems, such as the interferential filter.

**Fig. 4.** Photo of the Teletact 1.

The Lasercane N-2000[7] (fig. 5) is quite similar to the Teletact 1. The major difference lies in the fact that the Lasercane is integrated into a traditional white cane, whereas the Teletact 1 is a hand-held device. Thus, the cane can be used as usual, with the electronic aid switched off.

However, the maximum range of the Teletact 1 is greater: 10 m instead of 3.65 m for the Lasercane N-2000.

**Fig. 5.** Photo of the Lasercane N-2000. Photo © by Nurion-Raycal.

**Discussion** Infrared or ultrasonic devices are more easy to use than laser devices because they have a wide angle magnitude: users just have to point the device in front of them. Conversely, laser devices have a very narrow beam, so users have to scan their environment from side to side, which needs some training and good proprioception capabilities.

There is another issue with laser sensors: their 670 nm laser beam does not generally detect clean windows because it traverses glass without being reflected. Instead, objects located behind windows are detected.

However some users take advantage of this, for example to scan underground stations from train windows. Knowing that windows cannot be detected generally

[7] The Lasercane N-2000 was developed by the company Nurion-Raycal, see `http://www.nurion.net/`.

makes people feel uncomfortable. Black cars with metallic finish pose a similar problem: at high incidences, they do not emit a retro-diffused signal.

However, laser sensors are far more precise than infrared or ultrasonic ones, which gives them an invaluable advantage. That is why the designers of the Teletact 1 have released a new version of their device in collaboration with the LIMSI[8], which, among other things, embeds an infrared proximeter in addition to the original laser telemeter.

**Fig. 6.** Photo of the Teletact 2.

The so-called Teletact 2 (fig. 6) is equipped with a strong infrared superluminescent 950 nm proximeter [4, 5]. In case of both proximeter and laser telemetric detection, the system transmits telemeter information. When it senses the proximeter signal only, it sends a "window warning" signal to the user, in order to warn her that she may be approaching a translucent window. The proximeter works within a range of 3 meters, and gives a secured window pane and black car detection up to two meters.

The efficiency of the proximeter is due to the 950 nm wavelength and to the wider angle of the beam (about 20 ˚ ). Indeed, within a 20 ˚ angle at two meters, there is generally some "co-operative obstacle" that can be detected. For instance, when scanning a car, the device detects at least the handles of the doors.

### 2.2 Information Presentation

Once the distance information has been measured by the device, it must be conveyed to the user in some form. Existing devices use either audio or tactile interfaces.

**Audio Interfaces** The most simple kind of audio interface is implemented by the Lasercane N-2000: it simply uses a loudspeaker that emits a warning

[8] Laboratoire d'Informatique pour la Mécanique et les Sciences de l'Ingénieur, Orsay, France.

sound when there is an obstacle ahead of the user. As the device also has a tactile interface (see below), audio tones can be disabled in places where sound is restricted (e.g. museums, theaters, etc.)

Thus, the system provides a "go" – "no-go" information about the presence or absence of obstacles. If there is a signal, the traveler should stop. Conversely, in the absence of signal, the traveler may go on.

To convey more detailed distance information, the Teletact 1 uses 28 different musical notes, which correspond to 28 intervals of distance (the higher the tone, the shorter the distance). The 28 intervals of distance are unequal: they are smaller at short distances, because more precision is required for very close obstacles.

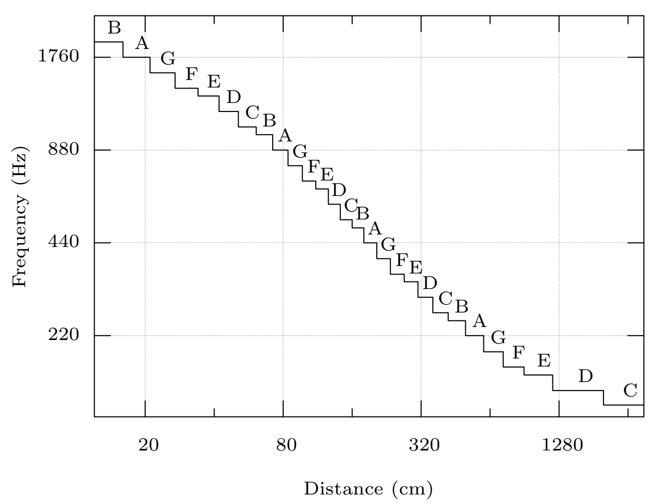


**Fig. 7.** Correspondence between musical notes and distance intervals (logarithmic scale).

The telemeter itself has a precision of about 1% in distance measurement. Therefore, the number of available musical notes is the limiting factor regarding precision. The device uses flute sounds over four octaves of the major scale (ranging from 131 Hz to 2.1 kHz). So there are 28 useful notes, assigned to 28 unequal intervals of distance in the 10 cm – 30 cm range (see fig. 7).

To get a profile of the obstacles ahead, it is necessary to scan the environment from side to side. It is not important to identify the resulting notes, but to interpret the meaning of the "melody" generated by the environment profile.

Therefore, users must develop some melody interpretation skills beforehand, so as to be able to recognize the melodies associated with common patterns such as corridors, stairs, and so forth. For this purpose, special training sessions have been set up that teach new users how to use their Teletacts.

To avoid confusion between the sounds emitted by the Teletact and possible surrounding noises (for example, coming from a street, or from a crowd), the Teletact uses earphones whose volume level can be adjusted at any moment to keep the distance signal comfortably clear. In addition, this solution presents the advantage of not disturbing the environment.

Some people have difficulties to rapidly merge the audio information with the proprioceptive one (movement of the wrist) to deduce spatial data. This information merge has to become instinctive after learning, because if the person has to think about it, using the device becomes too difficult an intellectual process.

After the necessary training course, fifty blind persons have successfully tested the Teletact 1 and used it to detect obstacles, avoid holes, and to find their way between parallel walls, etc.

The Miniguide has similar characteristics. The audio version uses a tone in earphones to indicate the distance to the closest object: the higher the pitch of the tone, the closer the object. The distance resolution of the audio interface can be set to either 2 cm or 20 cm, which is useful for people who prefer simpler feedback. To find openings only, it is possible to reduce the sensitivity of the system. Finally, there are special *watchdog modes* with very low power consumption used to detect people coming within a 2 meters range.

**Tactile Interfaces** The audio interface has various defects. Making a difference between musical sounds and street sounds is not a problem, in particular thanks to the earphones and the adjustable audio level. But when the ambient noise level often fluctuates, you have to adjust your proper audio level, and efficient automatic control in this case is very difficult.

To avoid the difficulties inherent to audio interfaces and to shorten the training phase, most devices propose tactile interfaces, generally in addition to an audio interface.

To illustrate this, let us look at the new tactile interface featured by the Teletact 2. It uses vibrating devices located under the user's fingers. Experiments were conducted with two, four and eight vibrating devices, and the four-device solution turned out to be the most successful.

The principle of this method is very simple. Each finger (except the thumb) is in contact with one and only one vibrating device (fig. 8). Each vibrating device corresponds to a distance interval. If an obstacle is detected within one of the four distance intervals, then the corresponding vibrating device is activated.

The Lasercane N-2000's principle is similar, but it stimulates the index finger only.

The Miniguide uses a different approach. In its tactile version, distances are indicated by the speed of vibrations. Thus, distance information is conveyed by the *modulation* of vibrations, not by their *spatial repartition.*

**Discussion** Using vibrating interfaces seems easier, because perception is more direct and intuitive. That is why beginners generally prefer this. However, people

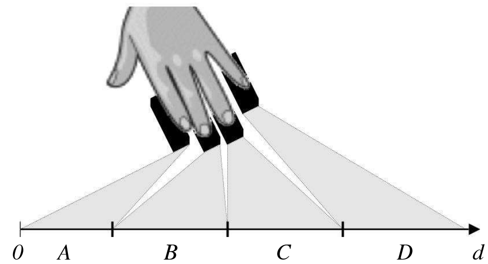


**Fig. 8.** The four vibrating devices correspond to four distance intervals ($A$, $B$, $C$ and $D$).

trained to audio output generally prefer this kind of interface and wish to stick to it because they find it more precise.

As regards the Teletact, the best results to date have been obtained by experienced users with the audio interface. In consequence, the best users tend to choose this solution. But the overall relative performance of the audio and tactile interfaces is not so different, in spite of a number of intervals being reduced from 32 to 4 in the tactile interface.

Furthermore, some advanced users of the tactile interface are beginning to perform tasks previously believed to be possible with the audio interface only (for instance, following a person in a crowd). Ongoing experiments will certainly tell us how these interfaces can be further improved.

However, what seems clear is that both interfaces need a good spatial representation, a good proprioception, and an active attitude: to scan the environment in order to look for obstacles. It is very difficult for blind people from birth to use it because of their lack of spatial representation abilities.

A long training period is required before being able to use a locomotion assistance device in everyday life, and it is especially difficult for elderly people (although there are some exceptions).

The main advantage of theses devices is the good anticipation users obtain and the optimization of navigation. It yields best results with very active blind people who have seen before, with ages ranging from 20 to 40 years old.

Table 1 summarizes the characteristics of the devices presented in this section.

## 3 Shortcomings of Current Systems

In the first section, we have presented a large variety of locomotion assistance devices for the blind. All these systems share a common characteristic: they only measure the distance to the closest obstacle, and convey it to the user. Their features seem to be somewhat limited when compared to what a leader dog can achieve. The gap is even greater between the mere distance information provided by these systems and the rich information and feedback permanently perceived by sighted people.

| Device | Measurement system | Modality | Resolution | Range | Angle magnitude | Autonomy | Learning Time | Size | Weight |
|---|---|---|---|---|---|---|---|---|---|
| **Miniguide** | ultrasonic | tactile and audio | 2 cm | 4 m | wide | 100 h | not estimated | 55 mm × 35 mm × 16 mm | 20 g |
| **Polaron** | ultrasonic | tactile and audio | binary (absence or presence of obstacles) | 1.2 m / 2.4 m / 4.8 m | wide | not known | not estimated | 27 mm × 50 mm × 162 mm | 257 g |
| **Ultracane** | ultrasonic | tactile | not known | a few meters | wide | 40 h to 60 h | short time | fits inside a cane | light |
| **Lasercane N-2000** | laser | tactile and audio | three levels | 3.65 m | very narrow | not known | not estimated | fits inside a cane | 450 g (with cane) |
| **Tom Pouce** | infrared | tactile | binary (absence or presence of obstacles) | 0.5 m / 1.5 m / 3 m | $20^\circ$ | 40 h | 3 to 4 months (10 to 20 sessions) | 55 mm × 75 mm × 17 mm | 90 g |
| **Teletact 1** | laser | audio | 28 levels | 10 m | $\ll 1^\circ$ | 1 h 30 min | 6 months (30 hours) | 200 mm × 100 mm × 50 mm | 450 g |
| **Teletact 2** | laser and infrared | tactile and audio | 28 levels or 4 levels † | 10 m | $\ll 1^\circ$ or $20^\circ$ ‡ | 15 h | 6 months (30 hours) | 170 mm × 50 mm × 20 mm | 180 g |

**Table 1.** Summary of the characteristics of existing devices. † 28 levels with the audio interface, or 4 levels with the tactile one; ‡ the angle magnitude is $\ll 1^\circ$ at long range (laser sensor), or $20^\circ$ at short range (infrared sensor).

This has led us to propose some new features that in our opinion would contribute to making much more usable devices. While these new features would greatly ease the traveling experience of blind people, all current systems lack them.

Thanks to locomotion assistance devices, blind people can avoid running into obstacles. This enables them to gain autonomy in crowded and dynamic places such as streets or train stations. However, people must know the premises beforehand in order to find their way. Indeed, existing devices cannot help people plan their path to some point, nor can they give them information about their location.

For this reason, blind people's autonomy is significantly reduced. When visiting a new place, they must come with a sighted person first, so as to learn how to find their way. Even if when they are sufficiently accustomed to a given place to easily find their way, they often have to ask sighted people to get any additional information. For instance, suppose that Bob is a blind person who regularly walks around an office building. Bob surely knows where the stairs are, so he basically can move alone. However, he surely does not know exactly where Alice's office is located if he has not yet visited her. Since he is blind, he cannot read maps nor signs, so he must resort to asking people at random.

Things can be even worse: if blind people are pushed, they may start to panic and be confused about where they are heading or even where they are located. In this case also they must ask somebody in order to resume walking in the right direction.

Car drivers are faced with the same kind of problem, at a different scale though. Indeed, apart from very well-known directions, it is hard to drive to a given place if one does not know the way. And in the same way as blind people sometimes no longer know their position and direction, car drivers sometimes find themselves driving on an unknown road, to an unknown destination.

As car drivers are concerned, the car industry has come up with a solution. Over the past few years, GPS[9]-enabled car navigation systems have become increasingly popular. They are no longer luxury accessories, but rather quite standard equipment.

A similar solution, adapted to the needs of blind people, would be even more helpful than car navigation systems are for drivers. It would prevent blind people from being completely lost, and guide them to their destination. With a navigation system, blind people could visit new places on their own, with no need of being accompanied by a sighted person.

Even when blind people know exactly where they are, they can be faced with another problem: they cannot easily get information about the surrounding objects. For instance, when in front of a door, they can have no clue as to what is located behind if they do not know beforehand.

In contrast, sighted people can easily get this kind of information: for example, by reading the door sign, they can learn the purpose of the room behind. For instance, an office door generally bears the occupants' names, which is very

[9] Global Positioning System.

important information. The nature of information can even be dynamic: for instance, if an occupant of an office unexpectedly leaves, she can leave a note on the door, stating the reason and the duration of her absence. Or, in the case of translucent doors, a sighted person can know if there is someone in the office simply by noticing whether the lights are on or off.

Since blind people cannot access to this kind of information by themselves, it would be very valuable if an handheld device could give them such pieces of semantic *descriptive* information about places or objects. Where current obstacle detectors can merely tell their users that, says, "there is an obstacle three meters ahead", future devices should be able to add that "this obstacle is a door", and that "this door leads to the boss' office". Advanced systems could be even further precise as to whether the boss is present at the moment or not.

No device of this kind has been released yet. However, now that we have seen the objectives for future assistance systems for the blind, let us summarize the functional requirements for these systems.

A navigation system will need two things:

– a means to know the precise location of the user. There are various localization techniques that will be described in section 4.2,
– structural descriptions of the visited places. For instance, a building will need to be described in terms of rooms, corridors, doors, etc. Section 4.3 will present a framework for environment modeling.

Besides, when giving descriptive information about a place, an information system will have to retrieve *semantic* information associated with it. Section 4.4 will explain how it is possible to store and query semantic annotations.

Structural and semantic descriptions can be either stored on the device itself, or retrieved from an ambient network when needed. We will not detail ambient networks *per se*, since they have become, or at least will rapidly become commonplace. Indeed, different technologies exist, such as third-generation mobile phone networks or ubiquitous WiFi[10] connections inside organization buildings and even at home. We guess that in most places users will very soon have at their disposal *several* different technological solutions to access the network.

## 4 Requirements for Future Systems

### 4.1 Introduction: Example of a Next-Generation System

This sub-section introduces a current research project that aims at providing guidance and descriptive information. Although quite limited in regard to the objectives aforementioned, this system is a good starting point for exploring the needs and technological solutions for the design of more complete systems.

[10] Wireless Fidelity. This term informally refers to wireless Ethernet networks, complying with the standards IEEE 802.11a, b and g.

A team of French researchers has launched the RAMPE[11] project [6], which aims at providing an information system for blind people when they take the bus.

Users carry small palmtop devices. For the time being, these devices are simply PDAs[12], on which only three buttons are available to control the system. This prevents blind users from confusing among many different buttons.

While walking in the streets, when a user comes close to a bus stop, his PDA pronounces the name of the bus stop through speech synthesis. If the user is interested in the given bus stop, she pushes a button. Then, the PDA requests the bus stop to advertise its position: a loudspeaker situated atop the bus stop emits a sound that the user is normally able to locate.

When the user arrives at the boarding location, his PDA reads the names of the corresponding lines. If the user shows her interest for a given line, the PDA reads the name of the stops on this line.

Technically speaking, the system is based upon WiFi access points[13] located in the bus stops. The *context of use* of the handheld device is simply the set of bus stops nearby. The detection of the latter is straightforward: the system just has to get the list of accessible WiFi access points, which is done by performing a standard WiFi scan. Therefore, when the user selects a bus stop, it is tantamount to selecting a WiFi access point. The device only needs to connect to it and use the resulting network connection to retrieve information (e.g. the line maps) or send commands to the bus stop (e.g. the command to advertise its position through the loudspeakers).

This system is currently being tested in urban bus lines. It is likely to be highly beneficial for blind people, because they can move around the lines and take buses without having to rely on passers-by. Indeed, they can know where to wait for a bus, which bus to take, and when to get off. This will improve their autonomy.

However, this system cannot be used in any place, as it relies on a very specific infrastructure installed in the bus stops. Indeed, the system does not just use any available WiFi network: it requires the access points to be installed in the bus stops, because WiFi access points coverage is used as a bus stop proximity sensor. Thus this system, although very well suited for its purpose, cannot easily be transformed into a *general* purpose guide for blind people.

A general purpose guide might as well rely on the presence of a wireless network, because such networks are becoming commonplace. However, one such guide cannot reasonably require *constraints* on the top *topology* of the wireless network.

---

11 Référentiel d'assistance aux personnes Aveugles pour leur Mobilité dans les transports publics et des Pôles d'Échanges, see `http://www.esiee.fr/~rampe/`

12 Personal Digital Assistant.

13 A wireless network is built around so-called *access points* that route data packets and act as beacons for network discovery and identification. Each access point has a unique identifier.

In this section, we introduce the requirements of future systems, in terms of user localization, environmental and semantic modeling. We introduce the concepts from a general point of view, although we will sometimes give more detailed information because we have already worked on implementing such a system.

### 4.2 User Localization

**Introduction** Several techniques have been proposed to track user location. To compare them, [7] proposes several discriminating criteria :

- *geographic vs symbolic*: some systems perform *geographic positioning*: they provide geographic coordinates with respect to a given coordinate system. Conversely, other systems perform *symbolic positioning* : they provide references (e.g. identifiers) to known objects. For instance, if Alice is in her office, a positioning system would yield the coordinates "`45° 34' 26'' N, 1° 47' 52'' E`" in the former case, and the identifier "`alices_office`" in the latter case,
- *absolute vs relative*: when performing *absolute positioning*, positions are given with respect to a fixed coordinate system. Conversely, when performing *relative positioning*, positions of points are given with respect to each other,
- *locality*: calculations can be performed locally, on a device carried by the user (*local localization*), or at the level of the surrounding infrastructure (*global localization*). In the former case, each positioning device keeps track of its own current position, whereas in the latter case, the infrastructure keeps track of the position of all users,
- *accuracy* measures the typical error of the system. For instance, a given system can have an accuracy of 10 meters, which means that *most of the time*, the position it reports is within 10 meters of the real position
- *precision* quantifies how often the positions reported by a positioning system are within the system's accuracy. For instance, if the former system has a precision of 95 %, it means that reported positions are within 10 meters of the real positions *95 % of the time*,
- systems can be designed to operate at different *scales* : worldwide, countrywide, inside a given building, etc.,
- system *costs* greatly vary. System designers must reach a trade-off between cost and precision and accuracy requirements.

Let us now examine the main positioning techniques that have been proposed (triangulation, proximity sensing, scene analysis, dead reckoning), and their technological variations. In a later sub-section, we will see how it is possible to increase precision and accuracy by combining *several* of these techniques.

**Triangulation** To perform triangulation, a system must know the positions of known reference points. It can then either measure the distances to these points (the technique is then called *lateration*), or the angles to these points (the technique is then called *angulation*). Such systems try to calculate coordinates : they are *geographic positioning* tools. Diverse techniques can be used to obtain the measurements.

The most common way is to install a radio-frequency beacon on each reference point. This way, users carry a positioning device (*local localization*) which measures distances by estimating the transmission time of travel, or by measuring the attenuation of the signal, which can then be used to deduce the position of the positioning device.

GPS is a well-known system that relies on triangulation. The reference points are satellites which move. The user receiver knows precisely the satellite trajectories, hence their position at any time. Satellites are precisely synchronized with each other and carry atomic clocks. The messages they send are time-stamped. User receivers measure *pseudo-distances*. That is the difference of the time-of-flight between the signals received from the satellites. This way, receiving the signals from three satellites suffices to calculate the two-dimensional absolute position. For three-dimensional positioning, a fourth satellite is needed.

GPS offers an accuracy of a few meters, especially since "Selective Availability[14]" has been turned off. However, GPS will work reliably only if at least four satellites are visible and have strong signals. This poses no problem in the open, hence the success of GPS among hikers and car drivers. Unfortunately, when users in cities are surrounded by high buildings the reception of satellite signals tends to be difficult. This phenomenon, called the *canyon effect* [8], often prevents reliable positioning by GPS in cities.

Inside buildings, the situation is even worse, and generally, GPS signals are very weak, and thus not usable. It is nevertheless possible to install a "constellation" of pseudo-GPS satellites (actually, beacons) inside a building. If the characteristics of pseudo-satellites are identical to these of actual GPS satellites, it is possible to use standard GPS receivers inside the building. This solution may be costly, because GPS signals are quite difficult to generate. Other solutions may be adopted, and are based on slightly modified GPS receivers [9].

It is also possible to use the same kind of triangulation techniques as used by GPS, but using other kinds of timestamped beacons, such as mobile phone base stations.

Triangulation-based location systems for mobile phones will be implemented shortly in the US, because the FCC[15] has passed rules for an Enhanced-911 service (E-911 Phase II) that will oblige telephone networks to be able to locate mobile phones within 50 to 300 meters when users call 911. Different techniques will be used. These include GPS (local localization), measurement of the angle of arrival (AOA) or time difference of arrival (TDOA) of mobile network beacons (local localization), and enhanced observed time difference (E-ODT). The latter

---

[14] Selective Availability was a "feature" of the GPS system that prevented people not accredited by the US Department of Defense from accessing the system's full accuracy. In practice, civilian users could only access deteriorated signals.

[15] Federal Communications Commission.

system requires additional beacons to be installed in mobile networks, but the positioning is ultimately performed locally by handsets.

Other systems are based on signal strength measurements in wireless networks, for instance mobile phone networks or WiFi networks. These systems can perform either local or global localization. This is because both fixed points (mobile phone base stations and WiFi access points) and mobile units all transmit signals. For instance, Microsoft's RADAR system achieves local localization by measuring the signal strength and signal-to-noise ratio of WiFi signals. However, these techniques remain imprecise due to uncertainties in radio propagation models [10] and signal strength measures.

All the solutions evoked earlier achieve absolute positioning. However, recent research work has been done in the field of UWB[16] communications [11]. In these systems, a vast quantity of users carry so-called UWB transceivers that communicate with each other and measure distances to one another. By exchanging their measures, the whole set of UWB transceivers can calculate their respective positions (*relative positioning* performed *at a global level*, due to algorithms distributed across the network of UWB transceivers).

**Proximity Sensing** The first commercial system of this kind was introduced by Olivetti in the early 1990s. It was called *ActiveBadge*. People had to wear badges that emitted their unique identifier through infrared signals every ten seconds. Infrared receptors were installed in each room, and constantly waited for badges to send their identity. This positioning system is symbolic and global. Location is done by room identifiers and is performed by the infrastructure.

A similar system has been introduced with Bill Schilit's ParcTab system [12, 13]: people carry PDAs that use an infrared-based computer network built around beacons installed in each room. By keeping track of the beacon currently in use to communicate with a given ParcTab device, the infrastructure can track the location of users.

In the CyberGuide system [14], infrared beacons are installed in rooms, and mobile devices perform local localization by identifying the nearest beacon.

Ultrasonic signals can be used in a similar way as infrared signals. For instance, the *ActiveBat* system, developed by AT&T (formerly Olivetti Research) in 1998, relies on ultrasound badges worn by people. Grids of ultrasound receivers are installed in the ceilings of rooms, thus providing a geographic location service by measuring the transmission time of ultrasonic signals. Conversely, in the *Cricket* system receivers are worn by users, and rooms are fitted with arrays of ultrasonic beacons. The main impediment to the massive adoption of these systems is their cost as many devices have to be installed in the environment for them to work.

Radio networks can also be used to sense the proximity of stations, although they do not try to perform triangulation as in the former sub-section. For instance, E-911 Phase I only uses the base station identifier as handset location.

[16] Ultra Wide Band.

Users are symbolically localized to a given cell of the mobile phone network. Similarly, a WiFi-enabled device can be localized by giving its access point identifier (ESSID, Extended Service Set Identification). Likewise, Bluetooth devices can "feel" the presence of each other by just scanning the Bluetooth devices located nearby.

RFID[17] is a technology that enables the detection of most generally passive tags. An antenna sends radio-frequency waves to the tags, which provides enough energy for them to answer by giving their unique identifier. This technology is used in such applications as identifying people when they enter buildings, tracking parcels when being routed for example [15].

Therefore, if an area is equipped with an RFID reader, it can detect people passing by, provided that (a) everybody wears an RFID tag, and (b) that the reader is sensitive enough. In this case, people positioning is mandatorily carried out by the infrastructure. Tags have to be installed in every "semantically-significant" place [16], and a server can then keep track of all users.

Conversely, it is possible to install tags at key locations in the environment and to equip user devices with tag readers. In this case, localization is performed at a local level. There is no need for the infrastructure to be computer-enabled; disseminating passive tags is sufficient. When passing by an RFID tag, a user is notified that she enters the corresponding area.

In these two examples, RFID tags are used to perform *symbolic* localization. A tag or a tag reader is in every place of interest, which allows identification. Studies have shown that it is possible to perform *geographic* localization as well by measuring the strength of signals received from the RFID tags [17].

Some issues arise with RFID technology, such as interferences between tags or between tags and metallic objects [18]. However, these problems can be overcome [19], so that RFID provides a reliable way of identifying objects by proximity sensing.

We have seen in subsection 4.2 that the strength of WiFi signals, with a propagation model, could be used to perform triangulation. However, as we have seen, this technique is unreliable. It is possible to use WiFi signal strength to perform localization without resorting to triangulation. At every point of a building, one can measure and record the strengths of the signals received from access points in a database. This way, a "fingerprint" is associated with each point. When someone wishes to know his position, he can measure the fingerprint of his present location, and look up in the database to find the closest recorded fingerprint, which gives the location.

This technique enables *geographic localization* through WiFi proximity sensing. This is not just symbolic localization with respect to the nearest access point. Commercial products of this kind have been developed by companies such as Skyhook Wireless[18] and Ekahau[19]. Intel's Placelab[20] uses a similar sys-

[17] Radio-Frequency Identification.
[18] `http://www.skyhookwireless.com/`.
[19] `http://www.ekahau.com`.
[20] `http://www.placelab.org/`.

tem, but it is applied to both WiFi access points, cellular phone beacons and Bluetooth devices, and not just to WiFi cells [20].

It is even possible to use FM radio stations as beacons, and this was demonstrated in Microsoft Research's RightSPOT system [21]. This system measures and ranks the signal strengths of several commercial FM radios. A database holds the probabilities of the possible rankings for each location. For a given ranking, a Bayesian classifier provides the most probable location. This system has limited accuracy: it can only precisely determine the suburb.

**Scene Analysis** These systems are based on video cameras and computer vision software. Cameras constantly take pictures of the rooms, and the processing power located in the infrastructure tries to identify people. A good example of theses systems is Microsoft Research's *Easy Living*.

People identification can be performed thanks to facial recognition. For instance it is possible to combine several modalities such as silhouette, skin color or face pattern. However, ambiguities can more easily be eliminated if people or objects have attached labels specially suited for computer vision, such as two-dimensional bar codes.

**Dead Reckoning and Map Matching** When no infrastructure such as GPS is available for position tracking, a device can independently try to keep track of its own position unaided. The basic idea is to embed an inertial unit that constantly measures the acceleration vector of the device, and then try to compute the device's position by integrating the acceleration twice. This method is called *dead reckoning*. Using this method, new positions are determined in a *relative* fashion, by estimating how far the user has moved since the last known absolute position. For example the last known position could be obtained from the GPS before it lost the signal.

Unfortunately, dead reckoning is very error-prone [22]. Computed positions are likely to deviate from the real positions because of cumulative errors due to the two successive integrations of the signals. To overcome this shortcoming, the *map-matching* method [23] suggests that the movements of people are restricted along well-defined paths on a map. Hence, it is possible to reduce deviation errors by computing the most probable position of the user *along a path* and not in every possible direction.

Map matching is used in car navigation systems, both in those relying on GPS as well as those relying on dead reckoning. This works well for car drivers because they are forced to follow roads, but in case of pedestrians movements are less predictable, especially in open spaces outside buildings.

Map matching is also extensively used in the field of robotics: even in robots that do not embed an inertial unit, the system precisely knows which actions the robot has made. Every moving and stirring command applied to the motors is kept in memory. Hence, the trajectory can be re-constructed and corrected by matching it against a map.

However, these techniques do not apply so well to human beings because of a number of factors.

1. when walking, human beings tend to not necessarily follow paths or roads, but rather walk around them. Especially in the case of blind people, users sometimes walk away from a “normal” path. As these are precisely the cases when assistance devices will be the most useful, they *must* keep track of people even when they adopt “abnormal” behaviors,
2. the human body is not as solid as a robot or car chassis. In consequence, if an inertial unit is attached to a human being, it will be tilted in every possible direction, its height will vary, and so on. All this is prone to large errors accumulating.

**Combination of Several Techniques** Each of the techniques presented in these sections has its own advantages and disadvantages. However, when several positioning techniques are available simultaneously, it is possible to increase the global accuracy and precision by combining their results.

For instance, if several geographic positioning techniques are simultaneously available, such as GPS data and WiFi triangulation the information can be fused. This can be done either by averaging them, or by using more advanced techniques such as Kalman filters.

Likewise, geographic localization and symbolic localization can be used in a complimentary manner. This way, it is possible to track the position of objects through computer vision (geographic localization), and to give a label to these objects when detecting their RFID tags or to a WiFi access point that they may carry [10].

### 4.3 Environment Modeling

To be able to locate users within their physical environment, a system must have a model of that environment. So far, we have worked on building descriptions only. However, the concepts introduced here can be extended to cover the full range of environment descriptions. This section explains how it is possible to represent the structure of buildings in a way that makes sense to a context-aware assistance device for the blind.

**Existing Solutions** Current formats for 3D scene description (such as VRML [24] or X3D [25]) focus on describing the mere visual *appearance* of environments. Indeed, they target *sighted people* and try to render scenes with as many visual details as possible. In contrast, we target *blind people* and thus we need to model the *structure* and *semantics* of environments.

The *structure* of architectural environments determines how architectural elements are organized to compose buildings. Structure is either visible (e.g. walls) or invisible (e.g. frontiers dividing a room in several zones, for instance smoking ones and non-smoking ones). These invisible frontiers are at least as important as the visible architectural elements. Therefore, they *must* be represented in the model, which 3D scene description formats do *not* take into account.

For data storage, geographical databases [26] represent an interesting framework. However, they do not allow *explicit* representation of structure, so we will not use them for building description.

**Modeling Building Structure** We have devised our own means to describe the structure of buildings. It allows descriptions to be built in a hierarchical fashion. Later we will see that such descriptions can be represented as XML[21] data. Our model is three-tier:

- *first tier:* we call *lexical elements* the simple (elementary) architectural elements, such as walls, doors, flights of stairs, and so on. Lexical elements are not necessarily *physical* elements. For instance, to delimit two zones in a room (say, a smoking one and a non-smoking one), we may define a *"virtual wall"* that is just like a wall, except that it has no physical existence;
- *second tier:* so-called *syntactic elements* are complex architectural elements, formed of several lexical elements. For instance, a room is defined by its walls, a stairway is defined by several flights of stairs and landings;
- *third tier:* syntactic elements are further aggregated in what we call *aggregation elements*. For instance, several offices can be gathered in a *cluster* called, say, "sales department".

A given description is composed of various elements that are instances of more general concepts. For instance, the laboratory's door is an instance of the general concept of a door. These concepts (or classes) take place in a concept hierarchy whose main branches correspond to the three families of objects.

All classes derive from an abstract common ancestor called an `Element`. This class has got three abstract subclasses, corresponding to the three tiers of our model: respectively `Lexical Element`, `Syntactic Element` and `Aggregation Element`. Concrete classes are then derived from one of these abstract classes, depending on what tier they belong to.

In a description, objects from the three tiers are bound together by two kinds of edges (cf. figure 9):

- *inclusion links* (solid lines) represent geometric inclusion between elements within a given tier;
- *composition links* (dotted lines) enable objects of tier $n$ to be composed of elements of tier $n-1$.

For instance, on figure 9, office number 2 is included in floor number 1. In turn, the later could be included in a given building.

As for composition, office number 2 is said to be composed by several walls (`w_1`, `w_2`, etc.). In other words, these walls *define* and *geometrically delimit* the office.

[21] eXtensible Markup Language.

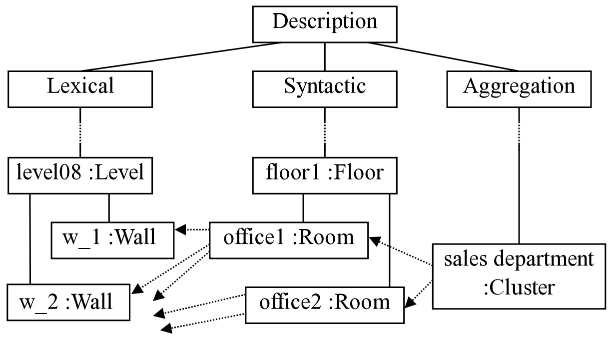


**Fig. 9.** Excerpt of an example description. The *sales department* is composed of two offices, which is in turn defined by some walls, and is included in a floor.

**Representation of Descriptions** This model has a strong hierarchical structure, so it is natural to use a representation format that makes this structure explicit. XML allows this type of explicit representation of structure, in addition to being widespread and universal. Therefore, we can use XML to store the three tiers of a description. For instance, the description of fig. 9 would lead to the following XML file:

```
<Description>
  <Lexical>
    ...
      <Level id="level-08" z="16m" height="2.20m">
        <Wall id="w-1" x1="2.3m" y1="4.2m" ... />
        <Wall id="w-2" x1="2.3m" y1="4.2m" ... />
        ...
        <Wall id="w-n" x1="12.3m" y1="13.2m" ... >
          <Door id="d-1" x="2.3m" width="1m" ... />
          <Window id="f-3" x="2.7m" y="0.5m" ... />
        </Wall>
      </Level>
    ...
  </Lexical>

  <Syntactic>
    ...
      <Floor id="floor-1">
        <Room id="office-1">
          <link ref="w-1" />
          <link ref="w-2" />
          ...
        </Room>
        <Room id="office-2" />
```

```
          ...
          <link ref="w-n" />
        </Room>
      ...
  </Syntactic>

  <Aggregation>
    ...
      <Cluster id="sales-department">
        <link ref="office-1" />
        <link ref="office-2" />
      </Cluster>
    ...
  </Aggregation>
</Description>
```

Inclusion links are represented implicitly through XML element imbrication, while composition links are represented explicitly thanks to XML `link` elements.

**Relevant Information** Using architectural descriptions like these, a locomotion assistance device will be able to determine where the user is. Then it will be able to provide the user with information located at the corresponding node.

However, we wish to give users the possibility to designate an object with a pointer, so that the system provides relevant information about the *pointed object* and not about the user's current location or nearest object.

Technically speaking, determining the pointed object is relatively easy if one already knows the *geographic* coordinates of the user. We can use a laser telemeter as a pointer. This way, with a compass and an inclinometer, one can calculate the trajectory of the laser beam, and thanks to telemeter data, one can finally deduce where the point of interest is.

Then, the whole issue boils down to determining the *level of detail*, i.e. the *granularity* of information needed by the user about the pointed object.

Indeed, too general information is useless, and too detailed information might not be understandable if the user does not know the associated context. To illustrate this, let us look at an example (see figure 10).

Suppose that a user is located on the second floor of the Computer Science (CS) Laboratory. She points through the window at an office on the first floor of another building, the Library building, located next to the CS building. What information shall we return? Information attached to the room, the floor, the building, the campus...?

If we provide the user with information about the campus, it will be far too general, and thus prove totally irrelevant, hence disturbing for her. Conversely, if we give her details about the very office at which she points at, (a) she will lack the context to understand this information, and (b) she surely is not interested

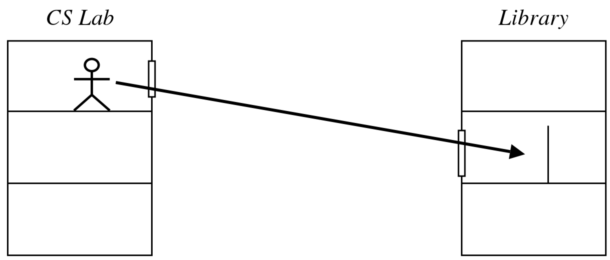


**Fig. 10.** A user points from one building to another one.

in an *office* located in another building anyway. In consequence, we must provide her with an intermediary level of information.

To solve this problem, we have devised an algorithm to determine a relevant level of information granularity.

We represent the scene as a tree (see figure 11): the user is located in $u$, and the pointed object in $p$. First, let us find the deepest node that is common to both the path leading from $u$ to the root, and the path leading from $p$ to the root. This node is labeled $c$ on Fig. 11.

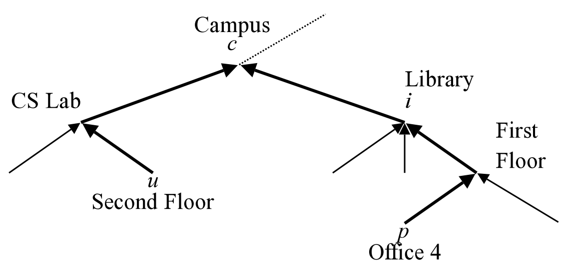


**Fig. 11.** A tree formally describing the situation depicted on figure 10.

What happens if we return information located at $c$ or above? Such information is too general, because it covers $p$ as well as $u$. Thus, it will probably be useless for the user, because being in the CS building they already know that they are on the campus.

In consequence, we should return information located *below* $c$. Therefore, this information will be on the sub-path leading from $p$ to $c$ since it must describe $p$. But it must be within the context of the user, because they could not understand it otherwise.

Therefore, we return information located in $i$, i.e. the most general information on the sub-path leading from $p$ to $c$. This corresponds to a *default granularity level*.

A locomotion assistance device may use this algorithm to determine the level of information granularity provided as a default in the interface. However, the user may wish to retrieve information located at another level. For this reason, the final user interface will offer some means of climbing up and down the tree, thus providing her with either more or less details about her object of interest.

### 4.4 Beyond Structure: Semantics

**Motivation** With the techniques described so far, one can design a handheld device that keeps track of the position of the user, and thus is able to locate her on a structural description of the environment. Moreover, when the user points in a given direction, we are able to determine a default object which is likely to correspond to a level of information granularity that makes sense to her.

The objects manipulated by the system (whether corresponding to the location of the user or to a designated direction) are structural elements (for instance, doors, rooms, etc.) A system can know these elements and their characteristics, and thus provide relevant information to the user: for example, "this is a door; and you have to push it to go beyond". For another door, the system may know that the user has to pull it open.

However, the user may be interested in getting additional *semantic* information. In the above example, the system would not only state that the user is standing before a door: it would add that this door leads to an office, and would then give out the office owner's name, the office function, and so on.

To do this, one may wish to associate semantic information to the structure description. More generally, such information can be used:

- to add *regulation* information to the structure, for example in order to tag restricted areas in a building,
- to identify objects, rooms, and zones;
- to represent connexity information;
- to add specific information to certain kinds of objects; for instance information about painters could be associated with sculptures in a museum.

**Solutions – Linking Semantics to Structure** Having stated the general problem above, we propose here our contributions and solutions. We introduce a particular means of representing semantic information, and linking it to the structure, using the model introduced in section 4.3.

We propose to represent semantic information using the Resource Description Framework (RDF) [27], an emerging W3C standard quite close to the theory of conceptual graphs [28]. We think that this choice makes sense because of the wide adoption of RDF by the industry. Thus, they are already a number applications able to handle RDF, and their number is steadily increasing.

In RDF, information is considered as a set of *assertions* that link objects of interest (sometimes called *resources*) together. Assertions are denoted by triples:

```
(subject, property, object)
```

An RDF description is composed of an arbitrary number of these triples, thus create multiple links between elements. This way, elements are connected in what is called an *RDF graph*.

Resources can be instances of classes. Several mechanisms have been proposed to describe class hierarchies. However, the *de facto* standard is OWL, which is the Web Ontology Language [29].

If we use RDF to represent semantic information, we end up with two worlds: one, a world of structure descriptions (possibly expressed in XML according to our model, described in section 4.3), and the other is a world of semantic annotations (expressed in RDF). These annotations exist only with respect to the underlying structure of buildings and they *need* to be connected to it.

To link semantic descriptions to the architectural elements of the structure description, we introduce a particular class of objects in the RDF descriptions. This is the class `Place`. Instances of `Place` correspond to architectural elements of the structure description and enable semantic graphs to be *anchored* in the structure. The actual correspondence is achieved through the use of a common identifier.

On figure 12, we can see the role played by the instances of class `Place`: these belong to the RDF graph, but they correspond one-to-one with the structural elements.

In practice, the structure in fig. 12 looks like this:

```
...
  <Building id="building-34">
    <Floor id="floor-1">
      <Office id="office-210">
        ...
      </Office>
    </Floor>
  </Building>
...
```

The triples of the associated semantic descriptions look like this:

```
(jack, hasCategory, crewMembers)
(crewMembers, hasAccessTo, privateAreas)
(jack, worksIn, office-210)
(floor-1, hasCategory, privateAreas)
...
(jack, rdf:type, Person)
(crewMembers, rdf:type, PersonelCategory)
(privateAreas, rdf:type, PlaceCategory)
(office-210, rdf:type, Place)
(floor-1, rdf:type, Place)
```

The structural elements `office-210` and `floor-1` have counterparts in the semantic description. The latter are instances of the class `Place`, and share the same identifier as their "structural" counterparts, which links the structure and semantics.

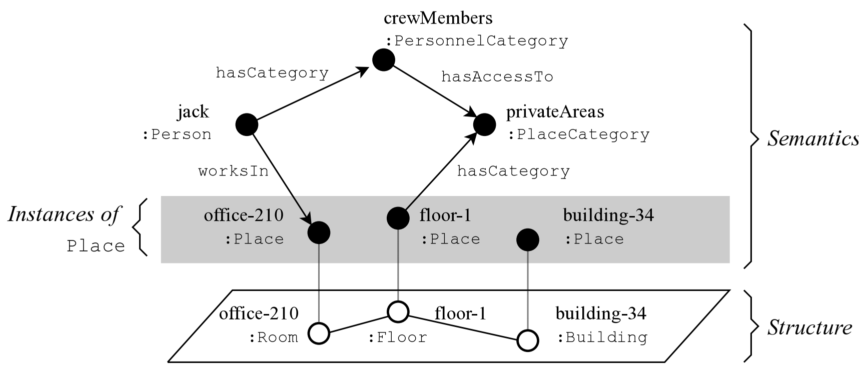


**Fig. 12.** Linking between semantics and structure.

**Ontology for Semantic Annotations** We have stated above that the objects of the RDF graphs are instances of classes and are defined using OWL. Here, we present some categories of objects that we think are important to include in an ontology for semantic annotations:

– *Locations:* in addition to the basic class `Place` already described, we propose a further two classes, `Function` that associates a function to a location (e.g. office, meeting room, etc.), and `PlaceCategory` that allows categories of places to be defined (with respect to an access restriction scheme, e.g. private area, public area, etc.);
– *People:* likewise, a class `Person` represents human beings, `Role` represents the roles played by people within an organization (e.g. employee, visitor, etc.), and `PersonnelCategory` allows categories of people to be specified. Again we use an access restriction scheme, staff member, intern, for example. It is not necessary to define a new class to represent persons. Instead, one can use the existing class `Person` from the FOAF (Friend Of A Friend) ontology [30];
– *Schedules:* to associate events to locations in a flexible way, we propose to define two classes: `Schedule` and `Event`.

The ontology may also define a whole set of RDF properties, used to specify relations between class instances. For example, `hasCategory` can be used to associate a `PersonnelCategory` to a `Person` or a `PlaceCategory` to a `Place`; `hasAccessTo` states that a `Person` or a `PersonnelCategory` have access to a `Place` or `PlaceCategory`.

**Reasoning** More than just enabling the description of complex relationships, a semantic description allows the definition of reasoning rules.

For instance, in the example of fig. 12 we have the following triples:

```
(jack, hasCategory, crewMembers)
(crewMembers, hasAccessTo, privateAreas)
(floor-1, hasCategory, privateAreas)
```

From that, it seems reasonable to deduce the additional triple:

```
(jack, hasAccessTo, floor-1)
```

Thanks to reasoning rules, it is possible to implement such *common sense* behaviors, and thus enable the automatic deduction of new triples from the set of existing ones.

## 5 Perspectives: Towards Ambient Intelligence

In 2001, Anind K. Dey defined the notion of *context* in human-computer interaction as *"any information that can be used to characterize the situation of entities [...] that are considered relevant to the interaction between a user and an application, including the user and the application themselves"* [31].

In this regard, the class of systems that we have presented here are *context-aware*. At the very least, their behavior, as well as the information they provide, depend on the location of the user. In reality, they could use many more aspects of context. For instance, taking ambient light into account provides useful information about activity in office buildings. For example when lights are off, employees are not likely to be at work.

Described as such, locomotion assistance devices would be applications for the numerous proposed frameworks for context-awareness, such as Anind K. Dey's own Context Toolkit [31] or Gaëtan Rey's Contextors [32]. They would be autonomous devices, constantly trying to extract contextual information from the surrounding environments, by reading physical sensors and querying data sources. This way, the environment is essentially unaware of the presence of blind users.

In contrast, we can imagine situations in which the environment is aware of the presence of users with special needs. For instance, if a door *knew* that a blind person is approaching, it could give hints through speech synthesis (*"Push to open!"*). Conversely, if a deaf person was approaching the door, it would not use speech synthesis, but rather flash a label with the word *"Push"* written on it.

In this case, information flows are much better than simple context-aware systems. Locomotion assistance devices behave as mobile *agents*, communicating with other agents embedded in the environment. This view corresponds to the notion of *ambient intelligence* [33] in which various objects have embedded computing capabilities as well as interaction facilities.

## 6 Conclusion

This chapter has described features, issues and solutions for next-generation mobility assistance devices for the blind. Figure 13 show an overview of what such a system could be, in terms of technological solutions.

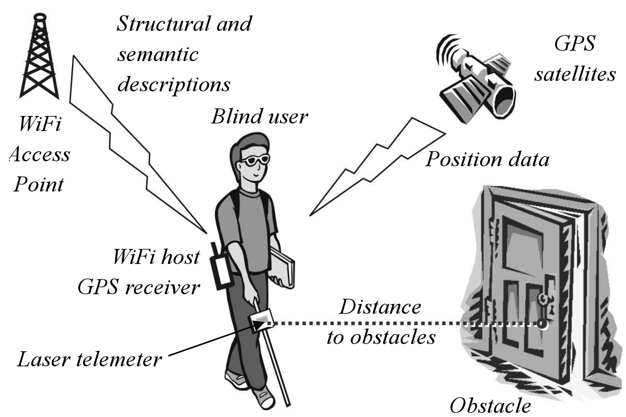


**Fig. 13.** Overview of a possible system.

This system builds upon existing electronic travel aids for the blind. It is based on a white cane, with a laser telemeter. However, it offers further features than just detecting close obstacles. Indeed, it permanently knows its position, thanks to a GPS receiver that works outside, and a WiFi positioning system that takes over the GPS inside buildings. Cartography and associated semantic descriptions are retrieved through the network, and provides guidance information as well as descriptive details about places.

This scheme remains relatively simple, because the handheld device communicates with its environment only to retrieve descriptions. However, a shift into ambient intelligence would turn the device into an *agent* communicating with various other agents located in the environment when the user visits new places.

The techniques needed by travel aids for the blind are not specific to this kind of applications. They can be applied to any context-aware mobile device, such as museum guide tours, travellers' assistants, and so forth.

## References


1. Jansson, G.: The functions of present and future electronic travel aids for visually impaired children and adults. In The Spanish National Organization of the Blind, ed.: Sixth International Mobility Conference, (1991) 59–64
2. Blash, B.B., Long, R.G.: Use or non-use of electronic travel aids in the united states. In The Spanish National Organization of the Blind, ed.: Sixth International Mobility Conference, (1991) 49–58
3. Farcy, R., Denise, B., Damaschini, R.: Triangulating laser profilometer as a navigational aid for the blind: optical aspects. Applied Optics **35** (1996) 1161–1166
4. Farcy, R., Bellik, Y.: Comparison of various interface modalities for a locomotion assistance device. In Miesenberger, K., Klaus, J., Zagler, W., eds.: Computers Helping People with Special Needs, Springer (2002) 421–428
5. Farcy, R., Bellik, Y.: Locomotion assistance for the blind. In Keates, S., Langdom, P., Clarkson, P., Robinson, P., eds.: Universal Access and Assistive Technology, Springer (2002) 277–284
6. Baudoin, G., Venard, O., Uzan, G., Paumier, A., Cesbron, J.: Le projet RAMPE: système interactif d'information auditive pour la mobilité des personnes aveugles dans les transports publics. In: Proceedings of Ubimob 2005, Grenoble, France (2005) 169–176
7. Hightower, J., Borriello, G.: Location systems for ubiquitous computing. IEEE Computer **34** (2001) 57–66
8. Chao, J., Chen, Y., Wu Chen, X.D., Li, Z., Wong, N., Yu, M.: An Experimental Investigation into the Performance of GPS-based Vehicle Positioning in Very Dense Urban Areas. Journal of Geospatial Engineering **3** (2001) 59–66
9. Samama, N., Vervisch-Picois, A., Francois, M.: La localisation en intérieur à l'aide de répéteurs GPS: vers un système de positionnement universel ? In: Proceedings of Ubimob 2005, Grenoble, France (2005) 21–28
10. Anne, M., Crowley, J.L., Devin, V., Privat, G.: Localisation intra-bâtiment multi-technologies: RFID, wifi, vision. In: Proceedings of Ubimob 2005, Grenoble, France (2005) 29–35
11. Daniele, N., Pezzin, M., Derivaz, S., Keignart, J., Rouzet, P.: Principle and motivations of UWB technology for high data rate WPAN applications. In: Proceedings of sOc 2003, Grenoble, France (2003) 84–87
12. Schilit, B.N., Adams, N., Gold, R., Tso, M., Want, R.: The PARCTAB mobile computing system. In: Proceedings Fourth Workshop on Workstation Operating Systems (WWOS-IV), IEEE (1993) 34–39
13. Schilit, B., Adams, N., Want, R.: Context-aware computing applications. In: IEEE Workshop on Mobile Computing Systems and Applications, Santa Cruz, CA, US (1994) 85–90
14. Long, S., Kooper, R., Abowd, G.D., Atkeson, C.G.: Rapid prototyping of mobile context-aware applications: The cyberguide case study. In: Mobile Computing and Networking. (1996) 97–107
15. Römer, K., Schoch, T., Mattern, F., Dübendorfer, T.: Smart identification frameworks for ubiquitous computing applications. Wirel. Netw. **10** (2004) 689–700
16. Ni, L.M., Liu, Y., Lau, Y.C., Patil, A.P.: Landmarc: indoor location sensing using active rfid. Wirel. Netw. **10** (2004) 701–710
17. Hähnel, D., Burgard, W., Fox, D., Fishkin, K., Philipose, M.: Mapping and localization with RFID technology. In: Proc. of the IEEE International Conference on Robotics and Automation (ICRA). (2004) 1015–1020

18. Floerkemeier, C., Lampe, M.: Issues with rfid usage in ubiquitous computing applications. In Ferscha, A., Mattern, F., eds.: Pervasive Computing: Second International Conference, PERVASIVE 2004. Number 3001 in LNCS, Linz/Vienna, Austria, Springer-Verlag (2004) 188–193
19. Brusey, J., Floerkemeier, C., Harrison, M., Fletcher, M.: Reasoning about uncertainty in location identification with rfid. Workshop on Reasoning with Uncertainty in Robotics at IJCAI (2003)
20. LaMarca, A., Chawathe, Y., Consolvo, S., Hightower, J., Smith, I., Scott, J., Sohn, T., Howard, J., Hughes, J., Potter, F., Tabert, J., Powledge, P., Borriello, G., Schilit, B.: Place lab: Device positioning using radio beacons in the wild. In: Proceedings of the 3rd International Conference on Pervasive Computing (Pervasive 2005). Volume 3468 of LNCS., Springer-Verlag (2005)
21. Krumm, J., Cermak, G., Horvitz, E.: Rightspot: A novel sense of location for a smart personal object. In Dey, A.K., Schmidt, A., McCarthy, J.F., eds.: UbiComp 2003: Ubiquitous Computing, 5th International Conference, Seattle, WA, USA, October 12-15, 2003, Proceedings. Volume 2864 of Lecture Notes in Computer Science., Springer (2003) 36–43
22. Fusiello, A., Caprile, B.: Synthesis of Indoor Maps in Presence of Uncertainty. Robotics and Autonomous Systems **22** (1997) 103–114
23. Kitazawa, K., Konishi, Y., Shibasaki, R.: A Method of Map Matching For Personal Positioning Systems. In: Proceedings of the Asian Conference on Remote Sensing ACRS 2000, (2000) 726–731
24. Web 3D Consortium: The Virtual Reality Modeling Language International Standard ISO/IEC 14772-1:1997. Technical report, Web 3D Consortium (1997) Available at `http://www.web3d.org/x3d/specifications/vrml/ISO-IEC-14772-IS-VRML97Wit%hAmendment1/`.
25. Web 3D Consortium: Extensible 3D (X3D) International Standard ISO/IEC FCD 19775:200x Final Committee Draft. Technical report, Web 3D Consortium (2003) Available at `http://www.web3d.org/x3d/specifications/x3d_specification.html`.
26. Hadzilacos, T., Tryfona, N.: An extended entity-relationship model for geographic applications. SIGMOD Record **26** (1997) 24–29
27. Manola, F., Miller, E., McBride, B.: RDF primer. Technical report, W3C World Wide Web Consortium (2003)
28. Sowa, J.F.: Conceptual graphs for a database interface. IBM Journal of Research and Development **20** (1976) 336–357
29. McGuinness, D.L., van Harmelen, F.: OWL web ontology language overview. Technical report, W3C World Wide Web Consortium (2003)
30. Dumbill, E.: Finding friends with XML and RDF. `http://www-106.ibm.com/developerworks/xml/library/x-foaf.html` (2002)
31. Dey, A.K., Salber, D., Abowd, G.: A conceptual framework and a toolkit for supporting the rapid prototyping of context-aware applications. Human Computer Interaction **16** (2001) 97–166
32. Coutaz, J., Rey, G.: Foundations for a theory of contextors. In: Computer-Aided Design of User Interfaces III, Proceedings of the Fourth International Conference on Computer-Aided Design of User Interfaces, Kluwer (2002) 13–34
33. Ducatel, K., Bogdanowicz, M., Scapolo, F., Leijten, J., Burgelman, J.C.: Scenarios for ambient intelligence in 2010. Final report, IST Advisory Group (ISTAG), European Commission (2001)