Towards an integrated approach to positioning

Christian Kray and Gerd Kortuem
Computing Department
Lancaster University
Lancaster, LA2 0PA, UK
{kray,kortuem}@comp.lancs.ac.uk

Abstract

The quality and availability of positional information (location, orientation) is a key resource that strongly influences the way in which many systems incorporating spatial information can provide services to their user. In this extended abstract, we discuss a number of issues and ideas on how to cope with low-quality positional information, how to improve that quality and how integrating qualitative and quantitative approaches may be beneficial in this context. We present a number of examples from systems that we developed in the past and that we are currently developing, and we point out possible improvements.

1 Positioning techniques

There is a large body of applications and systems that depend on knowledge about the location (and orientation) of their users and/or the position of the device used by them. These include (but are not limited to) mobile guides [3], car navigation systems [9], mobile phone based applications [8], geographic information systems [14] and ubiquitous computing environments [1]. There are various means for determining the current position of an entity. One of the most important ones is *device based positioning*, which is characterised by the fact that a (mobile) device itself determines its position - albeit in conjunction with an existing positioning or communication infrastructure. Device based positioning methods can be divided into three major categories:

- **Measurement based positioning**
  Measurement based positioning methods gather data from sensors and other positioning determination equipment and directly compute the device location. Perhaps the best known example of measurement based positioning is the Global Positioning System (GPS) (e. g. [11]). However, infrared beacons (e. g. [2]) and ultra-sound receivers (e. g. [1]) have also been used to the same effect. Furthermore, wireless communication networks such as GSM or 802.11 WLAN provide another means to determine the location of a mobile device (e. g. [3]).

- **Inference based positioning**
  Inference based positioning methods perform reasoning to improve the quality and/or precision of location information derived with measurement based methods. The most important example of inference based positioning is dead reckoning. Most inference based methods combine direct measurements with knowledge about past device locations and current movement patterns to infer a device's current location (cf. e. g. [10]). Inference based methods generate hypotheses: hypotheses are more or less reasonable but are not guaranteed to be correct.

- **Interactive positioning**
  Interactive positioning is a method that uses an interactive dialogue between system and user to determine the position [7]. A confirmation dialogue is the most simple example for interactive positioning: the system asks the user to confirm whether they really are at the position that was computed. A more soph-

isticated approach is to ask the user for an initial and rough estimate of the current location and to feed this information into an inference based method to improve accuracy.

Each of these three types of mobile positioning has advantages and disadvantages. Measurement based methods – by far the most popular approach to determine the location of an entity – can be very accurate but in some cases are not very reliable. Oftentimes, they do also depend on an external infrastructure (e. g. infrared beacons or satellites). Measurements and sensor information can exhibit systematic yet difficult to predict errors conditions due to the characteristics of the operating environment. For example, bad weather, narrow aisles and dense vegetation can severely deteriorate the reception of GPS signals. Infrared suffers from reflections and shielding, while ultrasound is prone to interference. Approaches based on network cells frequently face reception problems, e. g. in crowded places.

Inference based methods can work in situations where measurement based methods fail. In Deep Map [12], for example, we used a dead reckoning algorithm that takes into account contextual information such as previous locations, the current means of transportation, and the user's age and physical constitution to derive hypotheses about the user's speed, direction of movement and, ultimately, location. In this way, we were able to determine the user's location in situations in which measurement based methods alone would fail [5]. On the downside, it is oftentimes difficult to predict the accuracy of the derived location information. In addition, inference based methods can be computationally expensive and might not lend themselves to implementation on less powerful mobile devices.

Finally, interactive positioning is able to deliver location information even if no positioning determination equipment is available. For example, we have demonstrated that it is possible to accurately determine the user's current location by asking the user a small number of very specific questions regarding the visibility of landmarks – provided an accurate and detailed model of the environment is available [7]. However, interactive methods in general are intrusive and time consuming. Thus their application has to be carefully evaluated.

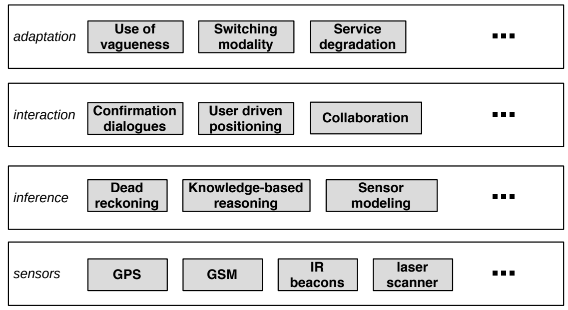

Figure 1: Different levels or layers involved in positioning: an overview.

2 Integrating positional information

The brief discussion of various techniques for positioning in the previous section illustrates that there is no single perfect solution. Instead, there are a number of approaches, which are well suited for certain scenarios. Consequently, it is worth investigating the integration of several techniques to overcome the disadvantages of each individual approach. This can happen on different levels. For example, we can merge information from a number of sensors (either several instances of the same type or different ones) to neutralise certain errors. This approach is frequently employed in cars or planes, which incorporate a large number of sensors such as GPS, compasses, or odometers. Alternatively, we can combine sensor-based approaches with inference-based techniques [13], or adapt the quality of a location-based service to the quality of currently available positional information. Figure 1 shows an overview over some techniques on the different levels, which can contribute to positioning.

The majority of approaches and systems computing the position of an entity rely on sensor data to do so. Consequently, there is a large body of techniques available on the sensor level including GPS, GSM, infrared beacons, and laser scanners. For virtually all of these approaches the raw sensor data (i. e. signal strength or time of flight) has to be processed in order to extract a position using algorithms such as particle filters or non-linear regression. Sometimes several sensors and/or positioning algorithms are combined to improve the accuracy, reliability or recency of positional information.

On the inference layer, an equally large number of complementary and competing approaches exist. These include reasoning about abstractions (such as topological relations), dead reckoning, and the explicit modelling of knowledge about the world or the sensors being used. They differ greatly in terms of resource consumption, expressional power and practical applicability. An example for a beneficial combination of techniques within this level is knowledge-based dead reckoning [5], where knowledge about the world and a user (e. g. the passability of a terrain or the user's means of transportation) influences the dead reckoning process.

The interaction layer encompasses various techniques that involve an interaction between the user and the system. While there are also a number of different approaches relying on interaction, their number is significantly smaller than the number of techniques available on the sensor and inference layer. Nevertheless, the set of mechanisms encompasses simple approaches such as confirmation dialogues (e. g. used to confirm that a position hypothesis is correct), user-driven positioning (e. g. users clicking on a map to localise themselves) and collaborative techniques (such as system-driven dialogues [7]). As these techniques all rely on some form of user interaction, it is straightforward to combine them. For example, a system-driven dialogue can be combined with a confirmation dialogue to ascertain the correctness of the position computed from the dialogue.

In addition to the three layers presented above, we can also include an adaptation layer, which is complementary to the other three. Here, we subsume all approaches that try to cope with positional information of varying quality or type. For example, instead of trying to improve on a lack of precision it is possible to explicitly use vagueness (e. g. by means of linguistic hedges such as "approximately" or "roughly", or by visualising the current position by a larger circle instead of a small point on a map). Further approaches include switching the modality in response to varying quality of positional information (e. g. from a map to speech output) or the (dynamic) degradation of the service depending on positional information (e. g. switching to a higher scale for a map display, or the use of induced frames of reference [6]). Combining different approaches on this layer is not as straightforward as it is within the other layers as adaptation on one dimension may preclude adaptation on another. For example, switching modality from a map display to spoken output precludes a change of scale on the map.

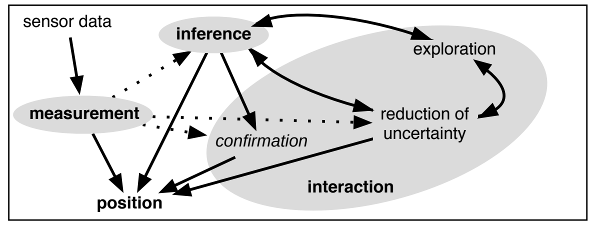

Figure 2: Processes involved in determining the user's position in Deep Map: measurement, inference, interaction.

Apart from combining different techniques within individual layers, it is certainly possible to combine a set of mechanisms from different layers. In the following two sub-sections, we present two examples which illustrate how such a combination can be achieved and what benefits can result from it.

2.1 Vertical integration: an example

Figure 2 illustrates how different positioning methods are linked within Deep Map's positioning service to provide for robust and adaptive positioning [5]. Deep Map is a mobile tourist guide for the city of Heidelberg. It provides its users with a number of services such as navigation support, information on sights and route planning. Its positioning subsystem uses measurements (sensor data) to directly compute the current position, and to provide its internal processes with an initial hypothesis. One of these processes is inference, which consists of a context-aware dead reckoning algorithm, and it is tightly linked to interactive processes. Whenever the result of the algorithm is not adequate for the task at hand – e. g. when it lacks precision – the output of dead reckoning is passed on to interaction for further processing. Similarly, when new information is acquired through interaction, the inference process is triggered in order to evaluate whether the user's current position can now be determined.

The combination of these methods allows for the dy-

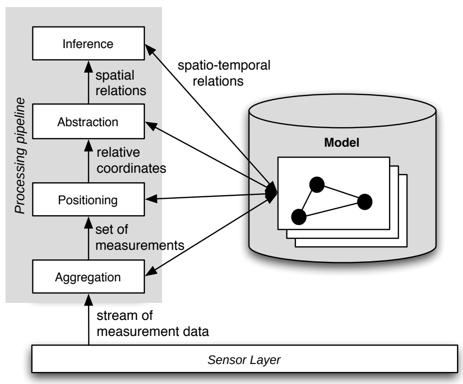

Figure 3: Processing pipeline in Relate: incremental determination of relative positions.

namic adaptation to the quality of sensor data. As long as the output of the sensors is sufficient to determine the user's position, no process except measurement is triggered. When the quality of sensor data degrades or when there is no data available, inference and/or interaction are triggered. It is even possible to determine the user's current position without any sensor data using the interaction process. It engages into a dialogue with the user, where it tries to compute the current location based on the visibility of objects in the environment [7]. However, depending on the quality of the underlying world model and the distinctiveness of the current environment, such a dialogue can be quite long and may not always yield a valid position hypothesis at the end.

2.2 Bidirectional integration: an example

The Relate project investigates relative positioning both on a technical (sensor) level and on a higher application level. Human interaction with artefacts on a table is one of the first scenarios that we focussed on, and which led to the realisation of the Relate dongle system [4]. We developed a small device equipped with ultrasound transducers that plugs into the USB port of a device, and that is capable of sensing its position relative to other dongles in its proximity. On the software side, the system contains a positioning subsystem depicted in Figure 3.

Raw sensor data (e. g. ultrasound signal levels) is gathered on the sensor layer, aggregated into chunks, and then transformed into relative coordinates. This transformation combines two different approaches: a non-linear regression algorithm and a simple map-tracing mechanism. The resulting coordinates serve as a basis for the inference stages in the pipeline (labelled 'abstraction' and 'inference' in the figure). Here, qualitative spatial relations are derived, which in turn inform the process of inferring further spatio-temporal relations. We are currently in the process of adding further mechanism to the inference layer, e. g. an approach that reasons about change from one configuration to another in order to assess the likelihood of a position hypothesis.

The bidirectional combination of positioning approaches in Relate's pipeline enables us to cope with a number of otherwise difficult situations. For example, when sensor data is scarce or when only few dongles are present the map-tracing algorithm performs better than the non-linear regression. The latter usually yields superior results in case of rich sensor data. The integration of inference techniques enables the system to better cope with noisy or false sensor readings that purely numerical approaches cannot identify easily.

3 Concluding remarks

The examples discussed in the previous section demonstrate that it can indeed be beneficial to combine different positioning techniques – not only on the same layer (sensors, inference, interaction, adaptation) but also vertically. It is possible to find combinations, where one technique can compensate the weakness of another one and vice versa. For example, qualitative approaches can help to assess the validity of positions computed by quantitative approaches. While we have pointed out a few beneficial combinations, there are further ones that still need to be explored. In addition, a systematic approach to adaptation is still missing. Further research is needed to formalise adaptation processes and to investigate how to best combine them.

Acknowledgements

The work presented here was funded by the EPSRC, grant GR/S77097/01 “Relative Positioning of Mobile objects” and by the Klaus Tschira Foundation, grants SISTO and SPACE.

References

[1] Mike Addlesee, Rupert Curwen, Steve Hodges, Joe Newman, Pete Steggles, Andy Ward, and Andy Hopper. Implementing a sentient computing system. *IEEE Computer*, 34(8):50–56, 2001.

[2] Jörg Baus, Antonio Krüger, and Wolgang Wahlster. A resource-adaptive mobile navigation system. In Yolanda Gil and David B. Leake, editors, *IUI 02 – 2002 International Conference on Intelligent User Interfaces*, pages 15–22, San Francisco, California, 2002. ACM Press.

[3] Keith Cheverst, Nigel Davies, Adrian Friday, and Christos Efstratiou. Developing a context-aware electronic tourist guide: Some issues and experiences. In *Proceedings of CHI 2000*, pages 17–24, Netherlands, 2000.

[4] Mike Hazas, Christian Kray, Hans Gellersen, Henoc Agbota, Gerd Kortuem, and Albert Krohn. A relative positioning system for spatial awareness of co-located mobile devices and users. In *Proceedings of the third international conference on mobile systems, applications, and services (MobiSys 2005)*, Seattle, WA, to appear.

[5] Christian Kray. *Situated Interaction on Spatial Topics*. PhD thesis, Computer Science Department, University of Saarland, Saarbrücken, Germany, 2003.

[6] Christian Kray. Inducing frames of reference. In *Workshop on Spatio-Temporal Reasoning at ECAI 2004*, Valencia, Spain, 2004.

[7] Christian Kray and Gerd Kortuem. Interactive positioning based on object visibility. In Stephen Brewster and Mark Dunlop, editors, *Proceedings of Mobile Human-Computer Interaction 2004*, pages 276–287, Berlin, Heidelberg, New York, 2004. Springer.

[8] Christian Kray, Katri Laakso, Christian Elting, and Volker Coors. Presenting route instructions on mobile devices. In W. Lewis Johnson, Elisabeth André, and John Domingue, editors, *Proceedings of IUI 03*, pages 117–124, Miami Beach, FL, 2003. ACM Press.

[9] Antonio Krüger, Andreas Butz, Christian Müller, Christoph Stahl, Rainer Wasinger, Karl-Ernst Steinberg, and Andreas Dirschl. The connected user interface: Realizing a personal situated navigation service. In Jean Vanderdonckt, Nuno Jardim Nunes, and Charles Rich, editors, *Proceedings of IUI 2004*, pages 161–168. ACM Press, 2004.

[10] Bu-Sung Lee, Wentong Cai, Stephen J. Turner, and L. Chen. Adaptive dead reckoning algorithms for distributed interactive simulation. *International Journal of Simulation*, 1(1–2):21–34, 2000.

[11] Sue Long, Rob Kooper, Gregory D. Abowd, and Christopher G. Atkeson. Rapid prototyping of mobile context-aware applications: The Cyberguide case study. In *Mobile Computing and Networking*, pages 97–107, 1996.

[12] Rainer Malaka and Alexander Zipf. Deep Map - challenging IT research in the framework of a tourist information system. In D. R. Fesenmaier, S. Klein, and D. Buhalis, editors, *Information and communication technologies in tourism 2000*, pages 15–27. Springer, Berlin, Heidelberg, New York, 2000.

[13] Reinhard Moratz and Jan Oliver Wallgrün. Spatial reasoning about relative orientation and distance for robot exploration. In Werner Kuhn, Michael F. Worboys, and Sabine Timpf, editors, *Proceedings of COSIT 2003*, pages 61–74, Berlin, Heidelberg, New York, 2003. Springer.

[14] S. Poslad, H. Laamanen, R. Malaka, A. Nick, P. Buckle, and A. Zipf. CRUMPET: Creation of user-friendly mobile services personalised for tourism. In *Proceeding of 3G 2001 - Second international conference on 3G mobile communication technologies*, pages 28–32, London, UK, 2001.