

A Semantics-based Framework for Context-Aware Services: Lessons Learned and Challenges

Theodore Patkos, Antonis Bikakis, Grigoris Antoniou, Maria Papadopouli,
Dimitris Plexousakis

Institute of Computer Science, FO.R.T.H.
Vassilika Vouton, P.O. Box 1385, GR 71110
Heraklion, Greece
{patkos, bikakis, antoniou, mgp, dp}@ics.forth.gr

Abstract. Location-based ubiquitous services have evolved significantly during the last years and are reaching a maturity phase, relying primarily on the experience gained and the utilization of recent technologies. This paper proposes a context-aware platform for mobile devices in dynamic environments, which uses Semantic Web technologies to model context information and advanced interactive map-based interfaces for accommodating pedestrians. To test and demonstrate the approach, a prototype has been developed and a number of further extensions are studied. Building and using this system has enabled us to identify the main challenges that need to be addressed for realizing the objectives of next generation semantics-based pervasive information systems.

1. Introduction

Location-based services and mobile ubiquitous applications represent an emerging research paradigm that aims at empowering users through personalized, context-aware and context-adaptive environments. Although mobile guides and navigation systems have come a long way since their first release, certain challenges have yet to be addressed, while new requirements are continuously becoming apparent. Thereby, new methodologies and advanced technologies are being developed that focus on providing flexible and extensible design models, proposing efficient solutions for dynamic environments, structuring solid context-aware infrastructures, as well as promoting openness and interoperability in response to the current highly heterogeneous reality.

Considering recent trends, the present paper proposes a framework for structuring advanced location-based and context-aware services that integrates up-to-date technologies and develops novel mechanisms and interactive interfaces by applying techniques and formalisms from the Semantic Web and Ubiquitous Computing domains. The objective is to explore the intelligent pedestrian navigation by implementing a Context Pedestrian Guiding platform (CG) for users in indoor environments. We focus on modelling and representing context using Semantic Web technologies for efficient processing and dissemination of context-based knowledge,

in order to develop services for mobile users. Integrating web ontology languages, such as RDF or OWL, with context-aware applications has manifold benefits. These languages offer enough representational capabilities to develop a formal context model that can be shared, reused, extended for the needs of specific domains, but also combined with data originating from other sources, such as the Web or other applications. Moreover, the development of the logic layer of the Semantic Web is resulting in rule languages that will enable reasoning with the user's needs and preferences and with the available ontology knowledge.

During the course of this project, we implemented a prototype that operates in the premises of the FORTH research facilities located in the technological and research park in Heraklion, supporting several of the basic functionalities that have been considered. Based on the experience gained during the development and usage of the platform, we reveal and comment on problems and challenges that must be considered in ubiquitous information systems engineering.

Our overall contribution is twofold. We exploit the expressive capabilities of Semantic Web technologies to develop a formal, shareable and extensible context model and integrate this model with a positioning mechanism and with flexible, interactive map-based interfaces to implement context-aware guiding services. Additionally, from a theoretical standpoint, motivated by our experiences, we raise and discuss numerous research issues and challenges that emerge, when attempting to understand the domain for enabling ubiquitous information systems to operate effectively.

The paper is organized as follows. In Section 2, related work is considered. The system architecture and implementation is presented next, while in Section 4 we describe the supported context-aware services. The paper concludes with a discussion on the main challenges that we encountered while developing and using the system.

2. Related Work

The existing implemented prototypes can be classified based on several criteria, namely the types of services that they provide (simple information guides, navigation systems), the context representation model (semantic vs. non-semantic), the environment they work in (indoor vs. outdoor), the type of the interface (map-based vs. text-based, interactive vs. non-interactive) and their adaptation capabilities. Two of the most prominent early mobile navigation systems were *Cyberguide* [1] and *GUIDE* [6], both of which were designed to provide visitors with up-to-date and context-aware information whilst exploring a city.

Only recently, systems that integrate Semantic Web technologies with navigation services have been developed. Two such systems are *mSpace Mobile* [17] and *COMPASS* [16]. They both use interactive map-based interfaces and aim to provide users with context-aware recommendations and map-based services based on the collection of distributed semantic information from the user's context. Furthermore, *mSpace Mobile* supports the lightweight semantic publication of comments, annotations and recommendations on any map element. *Navio* [13] and *OntoNav* [2] are two typical examples of navigation systems that use profile and context

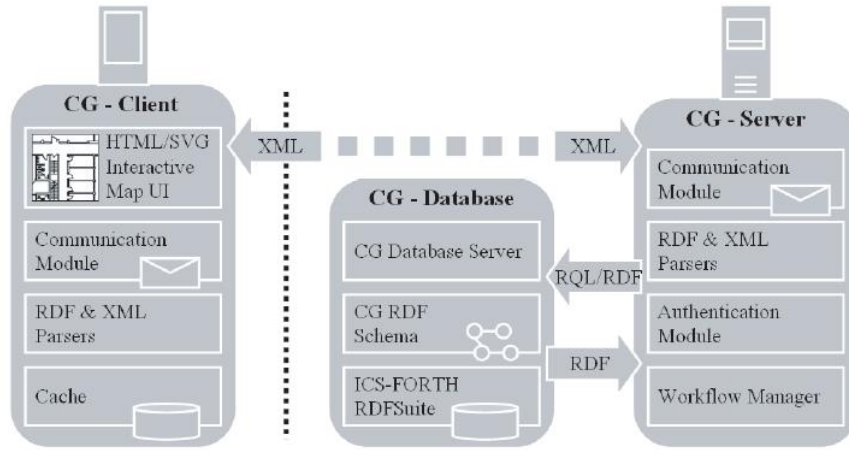


Fig. 1. The system architecture.

information for shaping routes and semantics-based methods for representing them. While the former proposes a pure ontology-based approach (a route modelling ontology), *OntoNav* uses a hybrid context model, which combines geometric with semantic modelling constructs.

Two are the distinct features of the Context Pedestrian Guiding platform, namely the extensible context ontology model and the interactive mechanism. The first includes representations of many aspects of context (such as, people, devices, locations, rooms), and combines spatial with temporal entities to model events. The second enables the users to not only issue queries about their context but also modify it, by updating the knowledge base through an interactive interface. Thus, the system can be used in two different ways: for automatically detecting changes in the user's context, but also as a medium to update the context in certain ways.

3. Context Pedestrian Guiding System Architecture

The platform's architecture, shown in Fig. 1, is designed to support user-centered and device-independent functionality, in order to provide the technical feasibility for building a multitude of context-based services. The framework, intended for indoor environments, is based on a centralized configuration and achieves a high level of transparency in inter-device communication, context management and service deployment.

Various context models have been proposed, but ontology-based are considered the most appropriate ones, as documented in [15], since they permit formal representation, richness of information and interoperability among different applications and devices. An RDF-based context model has been designed for use in the CG platform, aiming at addressing issues concerning semantic context

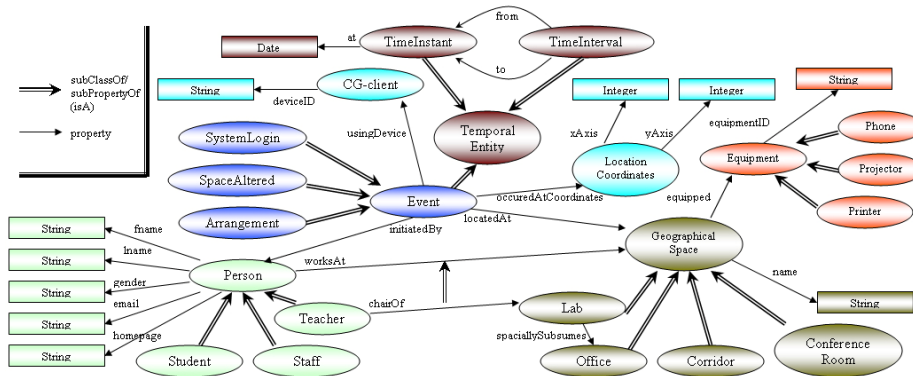


Fig. 2. The schema of the context ontology.
(different colours denote groupings of distinct concept clusters)

representation, knowledge sharing and context classification. The ontology captures the relevant notions of physical and semantic entities situated in the system and provides a formal vocabulary for structuring location and other context-aware services. The generic CG RDF Schema follows the abstract design principles of other large-scale pervasive computing ontology frameworks, such as the SOUPA project [5], and defines the high-level concepts of Person, Geographical Space, Location Coordinates, CG-Client and Event, as well as the relations among them. The model captures static information (i.e., device ID for CG-Client devices) and dynamic contextual data (events) that have both spatial and temporal extensions, such as user login, change of location in the building and other.

The CG Schema, along with all static and dynamic data, is stored by means of the ICS-FORTH RDFSuite [9]. RDFSuite is a suite of tools for RDF validation, storage and querying, using an object-relational DBMS. The language used for querying the database is RQL [10], a typed language following a functional approach, which supports generalized path expressions featuring variables on both labels for nodes (classes) and edges (properties).

A CG-Client module characterizes the devices with which a user can experience services provided by the CG platform. Different mobile devices can connect as CG-Clients in the system, such as PDAs or laptops, taking advantage of the flexible and portable Web-based model used for its development. CG-Clients connect with the CG-Server, a core component of our centralized architecture that acts as a middleware for processing, managing and transforming stored and inferred knowledge preserved in the CG-Database, while also offering the medium for synthesizing and publishing the desired services.

Implementation. For the CG platform to be adaptive to the complexity of context-aware scenarios, a number of up-to-date technologies and systems have been integrated that collaborate through a set of primitive methods (java api). An important component of location-based systems is the positioning mechanism that locates a mobile device or user in the environment, using some tracking technology. The CG prototype uses the IEEE802.11 infrastructure for both communication and

positioning. Specifically, it employs the Collaborative Location Sensing (CLS) [8] system that exploits the IEEE802.11 for positioning. The wide popularity of the IEEE802.11 network, low cost, and complexity advantages of using it for both communication and positioning make it an attractive choice. Typically, location-sensing systems that use an infrastructure of APs, employ a *signal-strength signature* or *map* of the physical space and contrast the run-time measurements with the map ([3], [11]).

The core functionalities of the CG-components, such as remote communication, multithreaded design and caching, have been implemented in Java. The CG-client user interface is composed of numerous Scalable Vector Graphics (SVG 1.2) files augmented by Javascript 1.2 functions for accessing the underlying Java procedures. This choice was not arbitrary. SVG is a vector graphics format suitable for the implementation of maps on the web, as well as on mobile devices [13]. These files are imported in HTML templates to allow familiarity with web-enabled interactions. Finally, all data is stored in RDF format using the ICS-FORTH RDFSuite semantic DBMS and exchanged using XML messages, while queries, expressed in RQL, are automatically generated, validated and manipulated.

4. Deployed Services

The system implements several services that enable users to obtain information about their context, to transparently utilize this information for supporting their everyday activity in their working environment and, also, to interact with the flexible map-based interfaces, by selecting or superimposing data objects on the map. Typical services that are included in the platform involve:

Semantic-based services. Users can have access to information stored in the platform's knowledge base regarding other users, such as their name, occupation, e-mail, homepage, current location in the institute, or regarding geographical spaces, such as room names, room equipment and other supported features. The client application automatically creates queries in a semantic language to retrieve static and dynamic data for performing the requested tasks. In addition, users can engage in interaction with entities appearing on their interface; for example they can send instant messages, publish semantically and contextually annotated notes, arrange presentations in conference rooms, check upcoming events etc.

Personalized services. Users that interact with a context-aware system are expected to experience independencies not only from the application and location perspectives, but also from the device perspective. Adapting content to the users' preferences and tailored to the capabilities of the device used is a priority for the CG system. We can use the user's profile information to support a variety of personalized services, such as message exchange among user groups, context sharing, group note creation etc. We also plan to augment the user's profile with privacy preferences that the system must take into account whenever it tries to retrieve a user's personal information.

Navigation / Route Finding. Using path planning algorithms, the context model, and the available information from the user's profile, we are developing services that help users navigate in the premises. These services will take into account not only the start and end point of the user's route, but also the user's special needs and characteristics, i.e., the authorization to pass through specific areas, potential corporal disabilities etc.

Calendar services. The integration of a calendaring application with the room reservation service is a valuable aspect of the CG system. We plan to employ a web calendar application that supports semantic descriptions of the calendar entries, in order to combine these entries with the room reservation semantic descriptions, but also with data originating from other semantic-based applications and sources.

A more detailed description of the supported services, along with a use-case scenario that demonstrates the ways that the system can be used as a valuable tool for the accomplishment of various everyday activities performed by people working in FO.R.T.H. are available in [4].

5. Challenges in Ubiquitous and Context-aware Implementations

Motivated by the experience and knowledge gained during the development of the CG platform, this section investigates certain problems that emerge and fundamental issues that must be addressed, during the implementation of ubiquitous, context-aware systems. We discuss the practical limitations and theoretical challenges of many fields of AI research, when applied to highly demanding realistic environments.

5.1 Distributed Planning and Multi-agent Coordination Techniques

Extending our platform with new services has made evident the importance of seamless collaboration between numerous devices (or respectively the agents representing them) that must work together to achieve common objectives. Even the more ordinary services like the management of a presentation in a meeting room may involve continuous cooperation between devices, as diverse as the room's projector, the lighting dimmer, the lecturer's laptop and many others. More complicated services require more sophisticated models of teamwork between devices, which differ in capabilities, characteristics and resource limitations. The establishment of common objectives among them entails a comprehensive understanding of the problem domain between participating entities and a desire to distribute their knowledge and to contribute their capabilities, in order to generate plans for achieving these objectives. Moreover, the assumption that the availability of devices is known beforehand is not valid in our platform, and, therefore, we have to come up with ways to arrange the formation of coalitions at execution time.

A sensor and device-rich environment along with software capable of fully leveraging their capabilities is, in fact, the domain of ubiquitous computing applications. If all the agents in the system were omniscient of the goals, actions and

interaction of their fellow community members and could also have unlimited computational resources, it would be possible to know exactly the present and future actions and intentions of each agent and systems could be perfectly combined. Obviously, this is not the case in most real-world systems and in ubiquitous IS, in particular. Despite the fact that our platform currently utilizes a centralized knowledge base, there are many situations where services would benefit by distributing reasoning and planning tasks among different devices. A typical example is when a user's connectivity with the network is undermined due to coverage range limitations or bandwidth restrictions, while he can still engage in ad hoc communication with other mobile devices in its vicinity, which can contribute resources and services. The self-organization and autonomy of devices are important design goals in several mobile computing environments, especially due to the very dynamic characteristics of the environment, frequent disconnections, and frequent interactions with other peer-devices. This decentralized self-organizing infrastructure is a nontrivial challenge for the realization of the ubiquitous computing vision.

Coordination does not imply either cooperation or reciprocation. deJardins et al. in [7] recognize two directions of research in approaching the problem of coordinating the actions of multiple agents in a distributed environment. *Cooperative Distributed Planning* is carried out by agents that have been endowed with shared objectives and representations, with the purpose to jointly develop and execute a plan in a coherent and effective manner. On the other hand, *Negotiated Distributed Planning* concerns agents engaging in negotiations over planned activities not to form good collective plans, but to try to influence each other to accommodate their individual preferences and ensure their local objectives will be met when viewed in a global context. Both directions have attracted significant research interest in an attempt to develop algorithms that perform efficiently considering a variety of metrics and at different environmental configurations. In the vast majority of situations, the ubiquitous computing domain involves agents engaging in Cooperative Distributed Planning, since it is assumed that personal devices of different users and appliances in a smart space will be willing to collaborate in achieving common objectives, when a particular level of trustfulness has been established.

5.2 Complex, Dynamic and Uncertain Environments

The majority of services provided by our platform requires a certain degree of reasoning and plan management skills by the participating agents. For instance, an agent wishing to print a document deliberates on whether to use the slow inkjet printer located near the user or the faster high resolution laser printer of the next room. While trying to enhance our platform with more aspects of context, the complexity of the planning task becomes computationally intractable. Planning and reasoning about action are well studied branches of AI, but have not yet been fully integrated in the ubiquitous computing domain. When research on planning within the AI community established the *classical planning problem* for the development of techniques for agents to generate courses of action to reach a desirable world state, it adopted a number of simplifying assumptions to delimit the domain. Some of these assumptions were very restrictive; the planning agent was considered omniscient, its

actions deterministic, atomic and simultaneous, the environment static and the only source of change in the environment was the agent itself, while no other exogenous event occurring. Of course, these simplifications do not persist in realistic planning situations and must be relaxed or completely eliminated when adapting planners to ubiquitous information systems.

Therefore, in studying the ubiquitous computing domain under the planning context, it is important to understand the classes of problems that planning methods face and the challenges that these problems introduce. The ubiquitous computing environment is an open and highly dynamic environment. Mobile devices connect and disconnect to the network, contributing services with durations that vary according to their expected presence in the environment and the availability of their resources. In fact, even actions, goals and sensor observations have a temporal dimension, whose duration may only be partially known in advance. The assumption of complete world knowledge can no longer persist; agents do not know *a priori* all other entities that are present at a specific time instance nor can they communicate directly with all of them. As a consequence, they have limited perception to acquire knowledge about the world they live in and have to generate plans preserving a level of uncertainty on both the state of the world, the available actions to achieve certain state of affairs and the outcome of those actions. Even the fact of committing different agents to certain tasks cannot be guaranteed to persist, since agents might disconnect before plan generation completes or new and more beneficial opportunities might arise, underscoring the importance of monitoring plan execution and replanning in opportunistic ways. The nondeterministic nature of the environment is emphasized by the recognition that not only agents, but also exogenous events occur in unpredictable and concurrent manner, affecting the state of the world. These observations essentially sketch a very challenging dynamic and uncertain environment with increased planning complexity. It remains questionable whether current efforts can attack this problem and provide computationally efficient and scalable solutions under these conditions.

5.3 Plan Representation and Sharing

The overall implementation of our context-aware platform was structured around a number of observations that specify the basic motivations of the domain, such as the desire to provide context-aware services to users that operate on a variety of mobile devices. In order for our architecture to support this scheme, we seek ways to represent the capabilities of these devices so that services can be adapted according to them. Profiles of devices must be flexible enough to capture both the complex actions that they can perform, but also decompositions of them to more lower-level actions, so that planning agents can understand and combine these actions to distribute responsibilities during service execution. There are problems that need to be overcome for achieving this type of interactions in our framework, since we assume that the existence of visitors possessing unknown mobile devices is going to be a common situation for the system. We also wish to exploit all the appropriate aspects of contextual information and to integrate them in the planning process, but also to dynamically associate with plan evaluation other dimensions, such as resource preservation, importance of execution, trust and privacy issues.

In order for any ubiquitous computing framework to support the challenges discussed so far, we also need to seek ways to represent device and service profiles, plans and goals in a manner that is mutually understood and correctly interpreted by all participating entities. Instead of exclusively focussing on improving our planners, other facets of the problem should also come into consideration that can refine the planning task. Recent studies suggest that it is not enough to just build better planners, which can generate plans for dynamic, uncertain, multi-agent environments ([12]). Autonomous agents also need to recognize which planning problems and opportunities to consider in the first place. They need to be able to weight alternative incomplete plans and to decide among competing alternatives. They need to be able to function intelligently in a ubiquitous environment in a way that comports with the inherent bounds on their computational resources. They need to share a common plan representation or a common ontology for describing their plans, goals and actions. And, they need to be able to exchange freely their plans and describe their action behaviors, privacy policies and authenticity certificates, in terms of expressive planning languages.

In order to attain the goal of planning in partially known environments with incomplete information, we have to endow agents with cognitive capabilities. When agents weight alternative subplans to reach a certain state of affairs, it is often useful to interleave reasoning in the process. We indeed believe that the ubiquitous computing domain raises many opportunities for exploiting reasoning to influence the efficiency of planning and reduce the search space. The most important leverage is to exploit the context in which an agent is situated, affecting its plan-generation capabilities. Context awareness has become one of the hot topics in recent computing research and reasoning on context can be an important source of information. It will enable them to constraint and better prioritize their goals, leading to computationally more efficient planning solutions in real world conditions.

6 Conclusions

The Contextual Pedestrian Guiding framework is a first effort to exploit the representational capabilities of the Semantic Web formalisms for the needs of a context-aware guiding application for the Institute of Computer Science of FORTH in Crete. The development and operation of a system prototype, and the study for the integration of more complex ubiquitous services have revealed us new research challenges, which require more effective solutions from the field of AI. The need for distributed, continual, and context-aware planning in a highly dynamic and demanding ambient environment dictates abandoning some strong assumptions of current methods and techniques, and inquiring new models that will enable the implementation of more efficient ubiquitous information systems. The study and integration of such models are part of our ongoing work on designing the extensions for our system.

References

1. Abowd, D. A., Atkeson, C. G., Hong, J., Long, S., Kooper R., Pinkerton, M.: Cyperguide: A Mobile Context-Aware Tour Guide. *Wireless Networks*, 3(5) (1996) 421–433.
2. Anagnostopoulos, C., Tsetsos, V., Kikiras, P., and Hadjiefthymiades, S.: OntoNav: A Semantic Indoor Navigation System. In *Proceedings of the 1st Workshop on Semantics in Mobile Environments (SME), MDM'05, Cyprus, May (2005)*.
3. Bahl, P., Padmanabhan, V. N.: RADAR: An In-Building RF-based User Location and Tracking System. *IEEE Infocom'00, Tel Aviv, (2000)*.
4. Bikakis, A., Patkos, T., Antoniou, G., Papadopouli, M., Plexousakis, D.: A Semantic-based Framework for Context-Aware Pedestrian Guiding Services. *Second International Workshop on Semantic Web Technology For Ubiquitous and Mobile Applications, Riva del Garda, Trentino, Italy August 28, (2006)*.
5. Chen, H., Finin, T., Joshi, A.: The SOUPA Ontology for Pervasive Computing. In *Ontologies for Agents: Theory and Experiences*, Valentina Tamma, Stephen Cranefield, Tim Finin and Steven Willmott (Eds), Springer Verlag, June (2005).
6. Cheverst, K., Davies, N., Mitchell, K., Friday, A., Efstratiou, C.: Developing a Context-aware Electronic Tourist Guide: Some Issues and Experiences. In *Proc. of the 2000 Conf. on Human Factors in Computing Systems (CHI-00)*, ACM Press, (2000) 17–24.
7. desJardins, M., Durfee, E.H., Jr., C.L.O., Wolverson, M.: A Survey of Research in Distributed, Continual Planning. *AI Magazine* 20(4) (1999) 13-22
8. Fretzagias C., Papadopouli M.: Cooperative location-sensing for wireless networks. *Second IEEE International conference on Pervasive Computing and Communications 2004, Orlando, Florida, March 14-17, (2004)*.
9. ICS-FORTH, The ICS-Forth RDFSuite web site, <http://139.91.183.30:9090/RDF>, (2006)
10. Karvounarakis, G., Alexaki, S., Christophides, V., Plexousakis, D., Schol, M.: RQL: A Declarative Query Language for RDF. In *Proceedings of the Eleventh International World Wide Web Conference (WWW'02) Honolulu, Hawaii, USA, May 7-11, (2002)*.
11. Ladd, A. M., Bekris, K. E., Rudys, A., Marceau, G., Kavraki, L. E., Wallach, D. S.: Robotics-Based Location Sensing using Wireless Ethernet. *The Eighth ACM International Conference on Mobile Computing and Networking (MOBICOM)*, Atlanta, (2002).
12. Pollack, M., Horty, J.F.: There's More to Life than Making Plans. *The AI Magazine*, 20(4) (1999) 71-84
13. Reichenbacher, T.: SVG for adaptive visualisations in mobile situations. *SVG Open Conference, Zürich, (2002)*.
14. Retscher, G. and Thienelt, M.: NAVIO - A Navigation and Guidance Service for Pedestrians. In: *Papers presented at the GNSS Sydney 2004 Conference, December 6-8, Sydney, Australia, (2004)*.
15. Strang, T., Linnhoff-Popien, C.: A Context Modeling Survey. *Workshop on Advanced Context Modelling, Reasoning and Management as part of UbiComp 2004 - The Sixth International Conference on Ubiquitous Computing, Nottingham/England, (2004)*.
16. van Setten, M., Pokraev, S., Koolwaaij, J.: Context-aware Recommendations in the Mobile Tourist Application COMPASS. *Proc. 3rd International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems (AH 2004)*, W. Nejdl, P. De Bra (Eds.). 26-29 August 2004, Eindhoven, The Netherlands, *Lecture Notes in Computer Science*, Volume 3137, Springer-Verlag, (2004), 235-244.
17. Wilson, M. L., Russell, A., Smith, D.A., Owens, A., and Schraefel, M. C.: mSpace Mobile: A Mobile Application for the Semantic Web. In *Proceedings of End User Semantic Web Workshop, ISWC2005, Galway, Ireland, (2005)*.