Supporting Knowledge Work by Observing Paper-Based Activities on the Physical Desktop

Stefan Dellmuth[1], Heiko Maus[1], Andreas Dengel[1,2]
[1]KM Department, German Research Center for Artificial Intelligence (DFKI)
[2]CS Department, University of Kaiserslautern
Kaiserslautern, Germany
Stefan.Dellmuth@senacor.com, {firstname.lastname}@dfki.de

Abstract

Despite the progress in personal information management, the desks of knowledge workers have remained untouched, justified by the claim that the paperless office is imminent. This led to a huge gap between the electronic world and the physical world. This paper combines the Semantic Desktop, methods of desk organization, and user observation together in order to bridge this gap. It focuses on paper artifacts and makes use of how knowledge workers organize them and interact with them. Paper artifacts are automatically observed, identified, and integrated into a model of the physical world. Moreover, paper documents are introduced into the personal knowledge space. A history of present and past desk layouts enables knowledge workers to search for information that was once present on their desks.

1 Introduction

During the last decades, huge advancements were made in electronic personal information management. In the 1970s, electronic devices started to mimick ordinary desks by employing the *desktop metaphor* as well as the *paper paradigm* [25, 8]. These advancements in conjunction with progress in scanning technology and document analysis sparked another trend: the belief that the *paperless office* was imminent.

However, the paperless office has not become reality. Even today, knowledge workers often prefer paper documents to electronic documents. The reason for that is simple: paper supports some of their most important daily tasks better than electronic documents [23]. Paper still excels in terms of tangibility, flexibility, easy navigation, and free organization on the desk.

Nowadays, the *electronic world* provides powerful software solutions such as the Semantic Desktop, which enables knowledge workers to create and maintain an electronic representation of their personal knowledge space for personal information management. However, the *physical world*, i.e., the knowledge worker's desk, has remained largely unknown and was often brushed aside by the claim that the paperless office was imminent and paper would disappear in the near future. However, paper still remains ubiquitous and the amount of paper documents in offices is even increasing.

In summary, paper has remained indispensable in modern offices [23]. Thus, the main goal of this paper is to investigate how both "worlds" can be brought closer together and how they can support each other in a way that benefits both. In order to achieve this goal, the paper rests on a particular workplace setup which yields more possibilities than an ordinary desk with a personal computer. The primary research question is formulated as follows:

> *Given an ordinary desktop, observed by a digital camera, as physical world and a computer with a Semantic Desktop as electronic world, how can the electronic world support activities in the physical world, and vice versa?*

The focus of the paper lies on the drawbacks of previous solutions in this research area. It proposes a system called *PhysicalDesktop* that aims at the following goals:

- Extend the personal knowledge space with a model of the physical world where paper-based information is stored semantically.
- Automatically extract as much information as possible from paper documents without any cognitive effort required of the knowledge worker.
- Make use of the physical desktop layout to support knowledge work.
- Bridge the gap between the physical and the electronic world in *both* directions.

The paper is organized as follows: the next section introduces the scientific background, including desk organization and paper-based activities in knowledge work. The third section gives a short summary of the external systems on which *PhysicalDesktop* is based. Then we present related work followed by the explanation of the main concepts behind *PhysicalDesktop*. Finally, the conclusion sums up the achievements and a comprehensive outlook on future work is given.

2 Background

There are two topics that contribute prominently to the physical world, i.e., the knowledge worker's observable desk: *desk organization* and *paper-based activities*. An extensive analysis of the scientific background was conducted in [3].

2.1 Desk Organization

Many publications identified two kinds of paper artifacts on the desk: documents and piles. While documents represent pieces of information on the desk, piles convey more than one meaning. Malone states that piles offer the function of *reminding* [14]. Kidd supports Malone's results and highlights the connection between the layouts of paper artifacts on the desk (so-called *spatial desk layouts*) and task management respectively (task) context [9].

Bondarenko and Jannsen conclude that knowledge workers change their desk's layout to reflect changes in their attention towards paper documents [2]. They state that knowledge workers move documents closer to a physical area of high interest, the so-called *focus of work*, when the urgency of the corresponding task increases. According to them, the distance between the focus of work and a paper artifact is an applicable measure of urgency for the task the document belongs to.

2.2 Paper-Based Activities

In the book *"the myth of the paperless office"* [23], Sellen and Harper conducted an extensive study about how office workers spend their work time on activities and how paper is involved. According to them nearly half of all document-based activities are solely based on paper documents.

For the purpose of having a sound and concise model, Sellen and Harper's activities of paper-based knowledge work were condensed to five abstract activities:

- *Acquiring Information* aims at getting information out of paper documents. It mainly encompasses different ways of reading.
- *Processing Information* is an abstraction of all activities that involve several sources of known information in order to produce new pieces of information.
- *Composing Documents* comprehends all activities that result in new paper documents.
- *Annotating Documents* encompasses all activities that lead to handwritten notes on paper documents.
- *Organizing Documents* summarizes all activities that lead to a change in the spatial desk layout.

Processing Information is tightly interwoven with Acquiring Information and in extreme cases, knowledge workers switch between both activities in short intervals. This is a strong indication for preemptive information acquisition, to satisfy information needs which might suddenly occur at a later date. As composing and annotating documents is already tackled in other projects [10, 21], the corresponding activities are left out of further considerations.

The abstract activities also tackle the activities for personal information management (PIM) [24] *transferred to the physical world*:

- *Finding & Re-Finding* centers around searching for information that is probably scattered across multiple paper documents and the desk itself. Problems arise if paper documents are buried deep within a pile. Moreover, it is not clear at first glance if the document is still available somewhere on the desk. Thus, finding and re-finding is related to Processing Information and Organizing Documents.
- *Keeping* paper documents presents a physical source of information. Thus, the problem of persistence occurs, especially if paper documents are required to be stored alongside electronic ones without creating clutter. Keeping is somehow related to all of the abstract activities. It shows that persistence is a crucial cross-sectional element in supporting knowledge work.
- *Organizing & Maintaining* paper documents on the desk is characterized by intrinsic behavioral patterns which depend on each individual knowledge worker. Adapting to these patterns as well as organizing and maintaining paper documents alongside their electronic counterparts poses serious problems. Knowledge workers have to be able to choose their preferred representations. Furthermore, they have to be assisted in finding the physical representation. Thus, this activity is clearly related to Organizing Documents.

Besides the activities, it is also of high interest how knowledge workers *interact* with paper documents on their desks and how these interactions contribute to the spatial desk layout. Several publications agree that knowledge workers place documents on the desk, move them around and, at some point in time, remove them from the desk [14, 9, 2, 11]. As far as piles are concerned, knowledge workers create, move, and delete them dynamically, according to their current needs [14, 9]. The basic document operations on piles are additions and removals.

It is worth noting that observing removals with camera-based desk observation systems poses serious problems. As the paper artifact leaves the observation system's sphere of influence, the removal can be detected indirectly at best, by recognizing that an artifact that has previously been on the desk is missing. The consequence of this circumstance is that desk observation systems can't claim absolute knowledge about the spatial desk layout. Hence, they can only give suggestions about artifacts being on the desk, based on internal reasoning.

Several studies [2, 23] make clear that detecting those interactions is just the first step in supporting knowledge workers. Beside the type of interaction, the destination of document placements and movements is a very useful piece of information. Additionally, if the desk observation system is able to recognize gestures on paper documents, the location and type of these gestures is also very important information.

3 External Systems

PhysicalDesktop uses three crucial external systems: the Semantic Desktop, the Context Service and iDesk.

3.1 NEPOMUK – The Semantic Desktop

The NEPOMUK[1] Semantic Desktop [4] focuses on supporting knowledge workers which are receiving, interpreting, and structuring information in their personal knowledge space. In general, the Semantic Desktop aims at transferring the idea and the technologies of the Semantic Web to desktop computers [19], providing a semantic layer on the distributed desktop information sources which allows represent and relate concepts and information objects across applications, e.g., the address book entry from MS Outlook of `Heiko Maus` with his papers or homepage. The core of this semantic layer is the *Personal Information Model Ontology* (PIMO) which is defined as "a formal representation of parts of the user's mental model" [20]. Thus, the PIMO allows to represent persons, projects, organizations, topics, etc. and their relations, e.g., `Stefan Dellmuth hasTopic physical desktop observation`, with information objects such as emails, documents, or webpages. As the Semantic Desktop provides an API for accessing the PIMO, concepts and structures that are created within the Semantic Desktop appear in every application which understands the model. Thus, the PIMO is a cornerstone for data integration in the environment of semantic technologies.

[1]open source: `http://nepomuk.semanticdesktop.org/`

3.2 Context Service

The Context Service is the reference implementation of the *user context for knowledge work*.[2] This context can be roughly defined as the user's context with respect to his current knowledge work state [22]. The current knowledge state is described by a set of contextual elements that the user currently has to deal with. The reference implementation uses information items from the PIMO as contextual elements. Information items which are observed contribute to the creation and enrichment of *context threads*. A context thread stands for a goal which the user pursues. The information items that are *stimulated* in the context of such a thread represent the information which was accessed in the course of working towards this goal. Naturally, a knowledge worker pursues many goals. Hence, the identification of the current context thread and context thread switches is a crucial requirement for the Context Service. Finally, the information that is represented by a context thread is used to support knowledge work in various ways, e.g., by providing input for proactive information delivery (PID) (cf. [16]).

3.3 iDesk

iDesk[3] is a desk observation system that has been developed by the DFKI IUPR department [12]. It consists of a digital camera and a software application. Its main goals are to detect paper documents and recognize pointing gestures on paper documents that reside on the user's desk. The main advantage of iDesk compared to similar solutions is that it is applicable in dynamically changing environments, as it is robust against changes in lighting and background.

4 Related Work

The approaches to combine the strength of the physical and electronic world are manifold. One approach is to *enhance paper documents with electronic media*. Its aim is to integrate information from the electronic world into ordinary workplace setups while preserving a natural look-and-feel. An exemplary solution — PaperWindows [6] — uses a motion capturing and a projection system to project application windows on real sheets of paper. Interaction is

[2]`http://usercontext.opendfki.de/`
[3]`http://ipet.iupr.org/demos.html#iDesk`

possible with a special pen and a set of gestures. The disadvantages of this approach include obtrusiveness (originating from the hardware setup) and the sole support for annotating documents.

Another approach aims at *providing a paper user interface* as an additional input method for computers. Recent solutions, like Print-n-Link [18], use Anoto's Digital Pen and Paper[4] technology. Print-n-Link is able to define links between user-defined parts of electronic and physical resources, such as paper and electronic documents, media resources and even source code. It also considers mobile scenarios very well. However, by relying on Anoto's technology, it loses the ability to perform well on arbitrary paper documents. This makes it impossible to interact with documents like (snail) mail, magazines, newspapers, or flyers.

A third approach is based on the idea of integrating physical and electronic documents into a unified solution. The main goal is to treat physical and electronic resources equally and thus combine the strengths of both worlds. Six years ago, the Paper Augmented Digital Document (PADD) paradigm was introduced. It centers on "digital documents that can be manipulated either on a computer screen or on paper" [5]. Thus, PADDs combine the strength of both worlds by an intelligent choice of representation. An exemplary solution — PapierCraft [13] — uses gestures to manipulate electronic documents by using paper documents as proxies. It extends the PADD infrastructure by introducing commands that users can trigger by using simple pen gestures. Another solution — Semantic eInk [21] — extracts semantic information out of handwritten annotations on paper documents and integrates them into the Semantic Desktop. Hence it makes paper documents, which present a new source of information, searchable, reusable, and shareable. In summary, the recent solutions to integrating physical and electronic documents excel at Annotating Documents. However, they neither support the other abstract paper-based activities, nor create links from the electronic to the physical world.

The last approach concentrates on making use of spatial desk layouts. Its goal is to keep track of paper documents on the desk and to provide links from the electronic to the physical world. An exemplary system was developed by Kim et al. [11]. It observes paper documents on a desk and provides links from the electronic world to the physical world. It is also worth noting that the system supports the full spectrum of spatial desk layouts, i.e., single documents and piles with according interactions. Its biggest disadvantage is the choice of a "scene graph"-like data structure to store the layout. This poses serious problems if the user observation misses interactions, e.g., when it was switched off on purpose.

Concluding, there exist some problems that none of the approaches and solutions tackle in a satisfying way. First, almost all of the solutions do not use suitable ways of persisting their data. Using the Semantic Desktop as persistent storage yields lots of new possibilities, e.g., connecting the acquired data with the knowledge worker's personal knowledge space. Second, most of the solutions can not cope with arbitrary or unknown documents. Moreover, they only support one or two of the abstract paper-based activities at best. Finally, all systems bridge the gap between the physical world and the electronic world from only one direction.

5 Concepts

Before the processing pipeline is explained, a few basic definitions have to be stated. First, the *physical desktop* is defined as the sum of all parts of the user's desk which are observed. The physical desktop in turn is defined as the sum of all *observed areas*. An observed area is a connected region on the desk which is observed by exactly one desk observation system.

5.1 Processing Paper Documents

In order to capture paper documents and transfer them to the electronic world, a processing pipeline was developed. First, a desk interaction is observed. The comprised paper document is identified. In case of an unknown document, the document is created in the PIMO of the Semantic Desktop. After that, the system is able to notify the user if the same paper document is already somewhere on the desk (to reduce redundant documents and clutter). If this is not the case, the history of interactions and the physical world model are updated. Finally, external services like the user's context are enriched or triggered. Figure 1 shows the resulting processing pipeline.

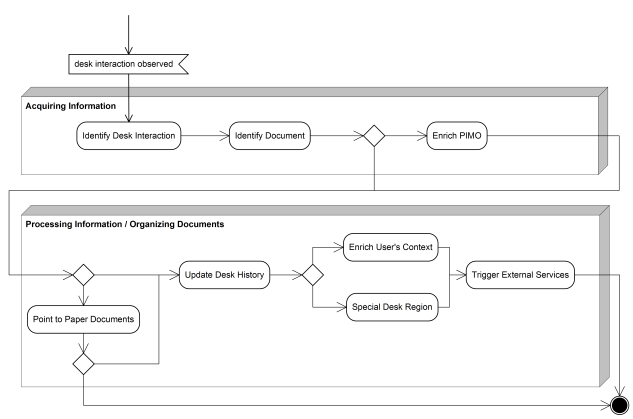

Figure 1. Processing Pipeline

The main task of the physical desktop observation is to capture the user's interactions with the desk and to send

[4]http://www.anoto.com

them in the form of desk operations, enriched with as much information as possible, to the pipeline. All operations are required to provide a location. The location information is composed of the observed area's unique identifier and, optionally, the position of the document's four corner points relative to the observed area. The documents themselves are represented as images. If a pile is involved, it has to be identified and the pile's unique identifier is provided with the desk operation. Finally, if a gesture was detected, the location of the gesture relative to the document and the type of the gesture are provided. Figure 2 shows the desk interaction ontology which was modeled using RDF/S.

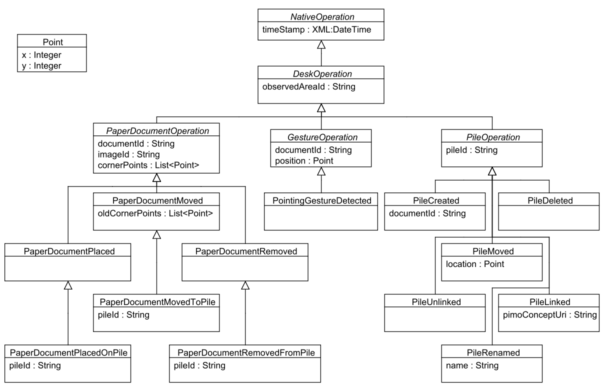

Figure 2. Desk Interactions Model

After the observation, the document has to be identified. It is crucial for the paper document identification to make as few errors as possible. Its quality has great influence on overall performance and user experience. Thus the identification algorithm is designed with robustness, accuracy, and error handling in mind. It uses either the Scale-Invariant Feature Transform (SIFT) or the Locally Likely Arrangement Hashing (LLAH) [17] algorithm. However, both algorithms merely state the best candidate from the database of known document images, along with a number of *votes*. As the number of votes which definitely indicates a known document is unknown and subject to change, an adaptive approach based on machine learning is used.

If the observed paper document is unknown, the PIMO is updated in the next step. For that purpose, the image of the unknown paper document is sent to an external OCR service. The result is then instantiated as a new document in the PIMO. Paper documents can also be linked to other information items to express an affiliation with a group of documents, a topic, a project, etc. Finally, they can be merged into existing documents in order to add the physical representation to existing electronic ones. This step in the pipeline also concludes the support for Acquiring Information. The remaining steps cover the activities of Processing Information as well as Organizing Documents.

If the observed paper document is known and already located somewhere on the desk, the system provides a link to its physical representations. The user is then able to choose whether to use one of the proposals or the recently observed document. This is useful to reduce clutter on the desk.

After this choice, the desk history is updated. The desk history is the central data store for all information concerning current and past paper artifacts and spatial desk layouts. It employs a model of the physical world which is explained in the next section.

Finally, the personal knowledge space is enriched further and external services can be triggered. The most important enrichment is that of the user's context. It allows to *stimulate* PIMO information items along with a *certainty* that the information contributes to the user's context. As all documents on the physical desktop are physical representations of PIMO information items, they can be easily stimulated. The estimate for the certainty takes into account the location of the document relative to the user (the farther away, the less important) and custom areas, the so-called *special desk regions*. Special desk regions are defined as parts of an observed area for which the normal rules of computing the certainty do not apply. Besides a custom certainty estimator, users are able to define custom rules that trigger actions upon certain interactions. With these possibilities, regions like "inboxes" (connect paper artifacts to a certain information item, e.g., "to read") and "outboxes" (remove paper artifacts from the physical desktop) can be defined. Moreover, this behavior can also be assigned to piles, thus enhancing their potential beyond passive paper piles.

Further external services are integrated such as a search service by using DynaQ[5]. Whenever a pointing gesture is detected, DynaQ is queried for similar documents to the one involved in the current interaction. The user can browse the results and use DynaQ to get detailed information and refine the query. Furthermore, services for easily annotating paper documents with PIMO concepts and transferring them into the Semantic Desktop such as the SCETagTool [15] or the follow-up iDocument [1] are also applied to support knowledge workers.

5.2 Physical World Model

As mentioned before, the desk history is based on a model of the physical world and stores all data which concerns current and past paper artifacts as well as spatial desk layouts. The model revolves around the paper artifacts and their *states*. While paper artifacts store all information that remains constant across all desk interactions, their states comprise all information that is subject to change. There are two kinds of paper artifacts: documents and piles. All paper artifacts have a unique identifier. Documents also contain an

[5] `http://dynaq.opendfki.de`

identifier which points to their image (which was taken by a desk observation system). Intentionally piles do not contain any additional information. This is due to the fact that piles have neither a constant representation on the physical desktop, nor in the PIMO. Thus, all information about them is stored in their state.

Desk artifact states serve several purposes. First, they pinpoint the location of a document or pile on the physical desktop by providing its source identifier and its location. The location is given as the paper artifact's *bounding area*. For documents, this is the smallest rectangular area that fully includes the documents. For piles, it is the union of the bounding areas of its documents. Second, they also pinpoint the location of a paper artifact in time. For that purpose, two timestamps are given, which mark the start and the end of the state. Finally, paper artifact states also record the type of interaction that created them as well as the previous and the next state.

Document states provide some additional information. First, they contain the exact location of the document, as recorded by the desk observation. If the desk interaction included a pointing gesture, the gesture point is also stored. Finally, the document's importance regarding the user's context is stored. Pile states contain all information that is known of a pile, i.e., the pile's name and a link to an information item, if one was established.

In order to counter limitations of desk observation system, the desk history provides a set of heuristics. They can be used to find out:

- whether a given location contains a paper artifact.
- if a document is assumed to be on the physical desktop.

The first heuristic is used to emulate most of the desk interactions that involve piles, as it is able to determine whether a paper document was placed on or moved to the top of another document or a pile. If either document deletions or document movements are not observable, the second heuristic is of particular use. By using a machine learning approach the heuristic decides whether a document is still on the physical desktop.

5.3 Graphical User Interface

The solution also features a graphical user interface for navigating the physical desktop. It consists of two applications: DeskViewer and DeskNavigator. While DeskViewer is the central interface for browsing current and past spatial desk layouts, DeskNavigator is used to get detailed information about paper artifacts.

Users are able to interact with the DeskViewer in several ways. For documents, the following actions are available:

- Show detailed information.
- Open its associated PIMO element with NEPOMUK.
- Associate the document with another PIMO element.
- Remove the document from the physical desktop.

While the former two actions deliver more information to users, the latter two provide fast and intuitive methods to correct the desk history. To perform the third action, users can simply drag-and-drop a PIMO element from NEPOMUK on the document's visualization.

DeskViewer also provides actions for piles:

- Show detailed information and browse the documents which belong to the pile.
- Open its linked PIMO element with NEPOMUK.
- Rename the pile.
- Link the pile to another PIMO element.
- Remove the pile from the physical desktop.

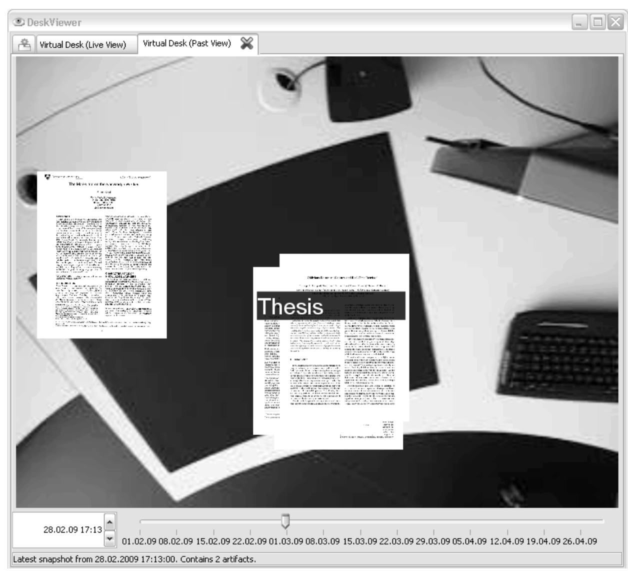

Figure 3. DeskViewer UI

Again, the first two actions deliver more information to users. The third and fourth action enable users to customize their user experience by naming the pile and linking a PIMO element with it, e.g., attaching a *topic* to the pile. In the case of linking, users can again drag-and-drop a PIMO element on the piles's visualization (see Figure 3).

Whenever users request detailed information of a paper artifact, the DeskNavigator appears. The application has a browser-like interface which consists of an "address bar" and tabbed views. The address bar can be used to search for paper artifacts by providing a part of their name or unique identifier. The tabbed views show a visualization of the paper artifact and a tree view which aggregates all past interactions to *time lines*. A time line is a set of desk interactions in chronological order which belong to a particular physical representation of the document. Selecting a leaf node opens a view of the physical desktop in the lower left corner. It shows the corresponding spatial desk layout and points to the location of the paper artifact.

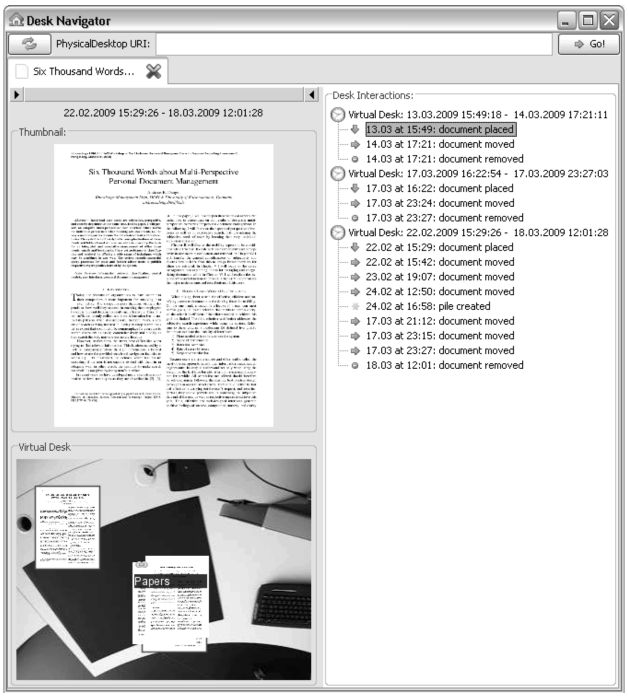

Figure 4. DeskNavigator UI

In summary, the DeskNavigator shows how paper artifact were represented and used over time. Figure 4 shows a screenshot of the application.

6 Conclusion

This paper addressed challenges in paper-based knowledge work. Based on the definition of the physical and the electronic world, it identified properties of paper documents and piles in desk organization and the set of desk interactions in paper-based knowledge work. The concept of *PhysicalDesktop* was then shaped based on the three abstract activities of Acquiring Information, Processing Information and Organizing Documents. Furthermore, a special focus was put on the ability to provide graded support and on unobtrusiveness.

PhysicalDesktop is integrated into the Semantic Desktop. It maintains a model of the physical world (i.e., the desk history) and enables paper documents to leave electronic traces in the personal knowledge space. By making use of spatial desk layouts, *PhysicalDesktop* enhances existing information stores like the PIMO and the user's context.

Furthermore, *PhysicalDesktop* employs an automatic information extraction process. A robust, learning document identification algorithm as well as external information extraction services aim at creating a rich electronic version out of paper documents which are observed on the physical desktop.

The most important result is that *PhysicalDesktop* bridges the gap between the physical and the electronic world from both sides. Paper documents are automatically transferred to the physical world model and the PIMO, which both are updated in real-time. At the same time knowledge workers retain full control over the model and its configuration. Information items can be used to point back to their physical representations. This allows for services like desk searches and turns paper documents into active objects of the personal knowledge space.

At the moment, *PhysicalDesktop* is available as an almost feature-complete research prototype. In order to conduct evaluations of the concept, an ideal emulation of the physical desktop, called *VirtualDesk*, was developed. Furthermore, a demonstrator was installed in the showroom of the DFKI (shown in Figure 5).

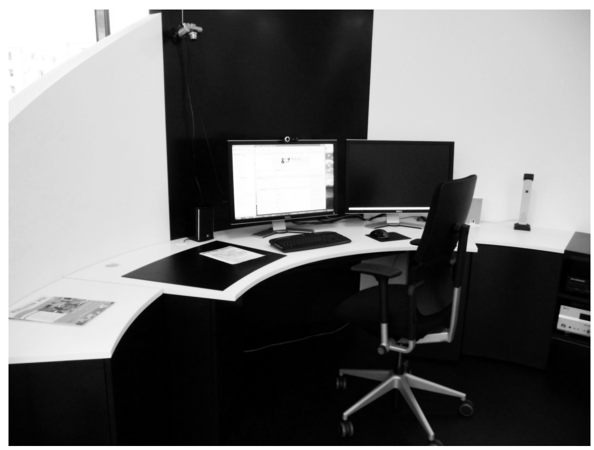

Figure 5. *PhysicalDesktop* **Demonstrator**

7 Future Work

As *PhysicalDesktop* is mainly a result of conceptual research, there is much room for improvement and enhancement. The bottleneck of the current implementation is the desk observation. Only a small subset of all proposed desk interactions and one gesture can be observed. Both the recognition of more desk interactions as well as gestures would lead to a substantial improvement. Another interesting thought is the integration of an eye-tracking device to recognize which paper documents the user looks at. This information could be used for real-time stimulation of paper documents in the user's context and for triggering real-time PID services.

Moreover, certain aspects of the solution are suitable to be supported by organizational structures. As most organizations have a central data store for electronic documents, the identification of paper documents and the retrieval of their electronic counterparts could be centralized. Furthermore, such data stores could be enhanced by multi function printers (MFPs) and automatic information extraction systems [7, 16]. By realizing these ideas, knowledge workers would be able to get rich electronic versions of paper documents by simply scanning them with an MFP and placing them on the physical desktop thereafter.

Finally, *PhysicalDesktop* comprises a basic form of error correction for the desk history. However, decisions by the desk heuristics and the document identification algorithm can't be corrected. A history of those decisions could be realized with the support to revert and change them.

Acknowledgments. We thank our colleague Christian Kofler from DFKI IUPR for his support with iDesk.

References

[1] B. Adrian, H. Maus, and A. Dengel. iDocument: Using Ontologies for Extracting Information from Text. In *5th Conference on Professional Knowledge Management*, 2009.

[2] O. Bondarenko and R. Janssen. Documents at Hand: Learning from Paper to Improve Digital Technologies. In *CHI '05: Proc. of the SIGCHI conference on Human factors in computing systems*, pages 121–130. ACM, 2005.

[3] S. Dellmuth. Supporting knowledge work by observing paper-based activities on the physical desktop. Master's thesis, University of Kaiserslautern, Germany, 2009.

[4] G. A. Grimnes, L. Sauermann, and A. Bernardi. The personal knowledge workbench of the nepomuk semantic desktop. In *The Semantic Web: Research and Applications. 6th European Semantic Web Conference, ESWC 09, Proceedings*, volume 5554 of *LNCS*, pages 836–840. Springer, 2009.

[5] F. Guimbretière. Paper augmented digital documents. In *UIST '03: Proc. of the 16th annual ACM symposium on User interface software and technology*, pages 51–60, 2003.

[6] D. Holman, R. Vertegaal, M. Altosaar, N. Troje, and D. Johns. Paper windows: interaction techniques for digital paper. In *CHI'05: SIGCHI conference on Human factors in computing systems*, pages 591–599. ACM, 2005.

[7] J. J. Hull and P. E. Hart. Toward zero-effort personal document management. *Computer*, 34(3):30–35, 2001.

[8] J. Johnson, T. L. Roberts, W. Verplank, D. C. Smith, C. H. Irby, M. Beard, and K. Mackey. The Xerox Star: A retrospective. *IEEE Computer*, pages 11–29, September 1989.

[9] A. Kidd. The marks are on the knowledge worker. In *CHI'94: Proc. of the SIGCHI conference on Human factors in computing systems*, pages 186–191. ACM, 1994.

[10] T. Kieninger and A. Dengel. The HCI Paradigm of Hyper-Printing. In *DAS '08: 8th IAPR Int. Workshop on Document Analysis Systems*, pages 591–598. IEEE, 2008.

[11] J. Kim, S. M. Seitz, and M. Agrawala. Video-based document tracking: unifying your physical and electronic desktops. In *UIST '04: Proc. of ACM symposium on User interface software and technology*, pages 99–107. ACM, 2004.

[12] C. H. Lampert, T. Braun, A. Ulges, D. Keysers, and T. M. Breuel. Oblivious document capture and real-time retrieval. In *1st Int'l Workshop on Camera Based Document Analysis and Recognition (CBDAR)*, pages 79–86, 2005.

[13] C. Liao, F. Guimbretière, K. Hinckley, and J. Hollan. Papiercraft: A gesture-based command system for interactive paper. *ACM Trans. Comput.-Hum. Interact.*, 14(4):1–27, 2008.

[14] T. W. Malone. How do people organize their desks? In *ACM Transactions on Office Information Systems*, volume 1, pages 99–112. ACM, 1983.

[15] H. Maus and A. Dengel. Semantic Annotation of Paper-Based Information. In *2nd Int'l Workshop on Camera-Based Document Analysis and Recognition (CBDAR)*, 2007.

[16] H. Maus, H. Holz, A. Bernardi, and O. Rostanin. Leveraging Passive Paper Piles to Active Objects in Personal Knowledge Spaces. In *Professional Knowledge Management. Third Biennial Conference, WM'05, Kaiserslautern, Germany*, LNAI, pages 50–59. Springer, 2005.

[17] T. Nakai, K. Kise, and M. Iwamura. Hashing with local combinations of feature points and its application to camera-based document image retrieval. *1st Int'l Workshop on Camera-Based Document Analysis and Recognition (CBDAR)*, pages 87–94, 2005.

[18] M. C. Norrie, B. Signer, and N. Weibel. Print-n-link: weaving the paper web. In *DocEng '06: Proc. of the ACM symposium on Document Engineering*, pages 34–43. ACM, 2006.

[19] L. Sauermann, A. Bernardi, and A. Dengel. Overview and Outlook on the Semantic Desktop. In *Proceedings of the 1st Workshop on The Semantic Desktop at the ISWC 2005 Conference*, pages 1–19. CEUR-WS, 2005.

[20] L. Sauermann, L. van Elst, and A. Dengel. PIMO - a Framework for Representing Personal Information Models. In *Proc. of I-Semantics' 07*, pages 270–277. JUCS, 2007.

[21] K. Schumacher, M. Liwicki, and A. Dengel. A paper-based technology for personal knowledge management. In *Proc. 5th Conference on Professional Knowledge Management*, volume P-145 of *LNI*, pages 289–298, 2009.

[22] S. Schwarz. A context model for personal knowledge management. In *Proc. of the 2nd Int'l Workshop of Modelling and Retrieval of Context (MRC 2005) in conjunction with IJCAII 2005*, pages 39–50, 2005.

[23] A. J. Sellen and R. H. Harper. *The Myth of the Paperless Office*. MIT Press, 2003.

[24] J. Teevan and W. Jones. The Disappearing Desktop: PIM 2008. In *CHI'08 extended abstracts on Human factors in computing systems*, pages 3917–3920. ACM, 2008.

[25] T. A. Wadlow. The Xerox Alto Computer. *Byte*, pages 58–68, 1981.