

RESEARCH OUTPUTS / RÉSULTATS DE RECHERCHE

A Graphical UIDL Editor for Multimodal Interaction Design Based on SMUIML

Dumas, Bruno; Signer, Beat; Lalanne, Denis

Publication date:
2011

Document Version
Peer reviewed version

[Link to publication](#)

Citation for published version (HARVARD):

Dumas, B, Signer, B & Lalanne, D 2011, 'A Graphical UIDL Editor for Multimodal Interaction Design Based on SMUIML', Paper presented at Workshop on Software Support for User Interface Description Language (UIDL '11), Lisbon, Portugal, 5/09/11.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

A Graphical UIDL Editor for Multimodal Interaction Design Based on SMUIML

Bruno Dumas
Vrije Universiteit Brussel
WISE Lab
Pleinlaan 2
1050 Brussels, Belgium
bruno.dumas@vub.ac.be

Beat Signer
Vrije Universiteit Brussel
WISE Lab
Pleinlaan 2
1050 Brussels, Belgium
beat.signer@vub.ac.be

Denis Lalanne
Université de Fribourg
DIVA Lab
Boulevard de Pérolles 90
1700 Fribourg, Switzerland
denis.lalanne@unifr.ch

ABSTRACT

We present the results of an investigation on software support for the SMUIML multimodal user interaction description language. In particular, we introduce a graphical UIDL editor for the creation of SMUIML scripts. The presented graphical editor is fully based on SMUIML for the underlying representation of data as well as for the dialogue modelling. Due to the event-centered nature of SMUIML, the representation of the multimodal dialogue modelling in the graphical SMUIML dialogue editor has been realised via a state machine. A real time graphical debugging tool is also integrated in the editor. Compared to existing multimodal dialogue editors, the SMUIML graphical editor offers a dual graphical and textual editing as well as a number of operators for the temporal combination of modalities.

Author Keywords

Multimodal Interaction, UIDL, Graphical Editor, SMUIML, HephaisTK

ACM Classification Keywords

H.5.2 Information Interfaces and Presentation: Graphical user interfaces, Prototyping, Theory and methods

INTRODUCTION

Multimodal interfaces aim to improve the communication between humans and machines by making use of concurrent communication channels or modalities. They have been shown to bring more comfort and better expressivity to users. Nevertheless, multimodal interfaces are difficult to realise due to a number of reasons. First, multimodal interfaces are typically composed of a number of state of the art recognition technologies, such as speech recognition or gesture identification systems based on pattern matching. Typically, developers have to master a number of these state of the art recognisers for different modalities in order to create advanced multimodal interfaces. Second, a combination of

input for the same modality can lead to ambiguous interpretations based on factors such as the ordering of input events, the delay between events, the context of use or specific user profiles. Fusion algorithms that take adaptation into account are therefore required. Last but not least, multimodal human-machine dialogue modelling is desirable in order to facilitate the development of complex multimodal interfaces.

A potential track of investigation to address the challenges introduced by multimodal interaction design is the use of a modelling language in combination with a multimodal framework and development environment. A multimodal user interface description language (UIDL) forms the key element in such an approach. The UIDL is used to define the behaviour of the multimodal framework, to perform the dialogue modelling and as the underlying format for the GUI development environment. A multimodal user interface description language is typically situated at the Abstract User Interface (AUI) layer. Furthermore, software support for the UIDL is provided for the definition, modelling or interpretation of user interfaces descriptions.

We present our explorations of such a language-based approach in the context of the Synchronized Multimodal User Interfaces Modelling Language (SMUIML) and the corresponding software support. In particular, we present a graphical UIDL editor for SMUIML and discuss its support for designing multimodal interactions. The graphical editor offers an alternative to the pure text-based editing of our XML-based language which is often more tedious and can easily lead to errors. We start by discussing related work in the context of modelling languages as well as graphical editors for multimodal interaction design. We then introduce the SMUIML language as well as some results from our research on the language design for modelling multimodal interaction. This is followed by a description of the different supportive software components for the SMUIML language with a particular focus on the graphical UIDL editor. After an overview of the planned future work, we provide some conclusions.

RELATED WORK

There have been a number of formal language approaches for the creation of multimodal interfaces over the last decade. Some of these approaches are positioned in the context of a

multimodal web, propagated by the World Wide Web Consortium's (W3C) Multimodal Interaction Activity and its proposed multimodal architecture¹. This theoretical framework describes the major components involved in multimodal interaction, as well as potential or existing markup languages to be used to relate those different components. Many elements described in this framework are of practical interest for multimodal HCI practitioners, such as the W3C EMMA markup language² or modality-focused languages including VoiceXML³, EmotionML or InkML⁴. The W3C framework inspired Katsurada et al. [9] for their work on the XISL XML language. XISL focuses on the synchronisation of multimodal input and output, as well as dialogue flow and transition. Araki et al. [1] propose the Multimodal Interaction Markup Language (MIML) for the definition of multimodal interactions. A key characteristic of MIML is its three-layered description of interaction, focusing on interaction, tasks and platform. Ladry et al. [11] use the Interactive Cooperative Objects (ICO) notation for the description of multimodal interaction. This approach is closely bound to a visual tool enabling the editing and simulation of interactive systems, while being able to monitor system operations at a low level. Stanculescu et al. [15] followed a transformational approach for the development of multimodal web user interfaces based on UsiXML. Four steps are necessary to go from a generic model to the final user interface. One of the main features of their work is a strong independence from the available input and output channels. A transformational approach is also used in Teresa XML by Paterno et al. [13]. Finally, at a higher level of modeling, NiMMiT [2] is a graphical notation associated to a language used for expressing and evaluating multimodal user interaction.

Graphical editors for the definition of multimodal dialogues can broadly be separated into two families. These two families differ in the way how a dialogue is represented which is often driven by the underlying architecture. On the one hand, stream-based architectures favour a direct representation of data streams, with building blocks consisting of processing algorithms that are applied to the streams in a sequential manner. There has been a trend for graphical editors for stream-based multimodal architectures in the past few years. Petshop for ICO [11], Squidy [10] or Skemmi [12] for Open-Interface are examples of graphical editors for stream-based architectures. On the other hand, event-driven architectures result in a state machine-based representation of the multimodal human-machine dialogue. In this category, fewer examples exist for the representation of multimodal interaction, the most prominent one being IMBuilder from Bourguet [4].

Note that the graphical editors introduced in this section have all been built from scratch and they are not based on a previously defined formal language, with Petshop for ICO forming the only exception.

¹<http://www.w3.org/TR/mmi-arch/>
²<http://www.w3.org/TR/emma/>
³<http://www.w3.org/TR/voicexml20/>
⁴<http://www.w3.org/TR/InkML/>

THE SMUIML LANGUAGE

SMUIML stands for *Synchronized Multimodal User Interaction Modeling Language*. As its name implies, SMUIML aims to offer developers a language for describing multimodal interaction, expressing in an easy-to-read and expressive way the used modalities, the recognisers associated with a given modality, the human-machine dialogue modelling, the various events associated to these dialogues and the way these different events can be temporally synchronised.

The SMUIML language is divided into three layers, with a different level of abstraction for each layer. Figure 1 highlights these three levels of abstraction. The lowest level details the different modalities which will be used in the context of an application, as well as the particular recognisers to be used to access the different modalities. The middle level addresses input and output events. Input events are called *triggers* and output events *actions*. Triggers are defined per modality which means that they are not directly bound to specific recognisers and they can express different ways to trigger a particular event. For example, a speech trigger can be defined in such a way that the words “clear”, “erase” and “delete” will all lead to the same event. Actions are the messages that the framework sends to the client application. The top level of abstraction in Figure 1 describes the actual human-machine dialogue by means of defining the contexts of use and interweaving the different input events and output messages between those contexts. The resulting description takes the form of a state machine similar to Bourguet's IMBuilder [4]. The combination of modalities is defined based on the CARE properties [6] as well as on the (non-)sequentiality of input triggers. As shown in Listing 1, these levels are directly reflected in the basic structure of the language.

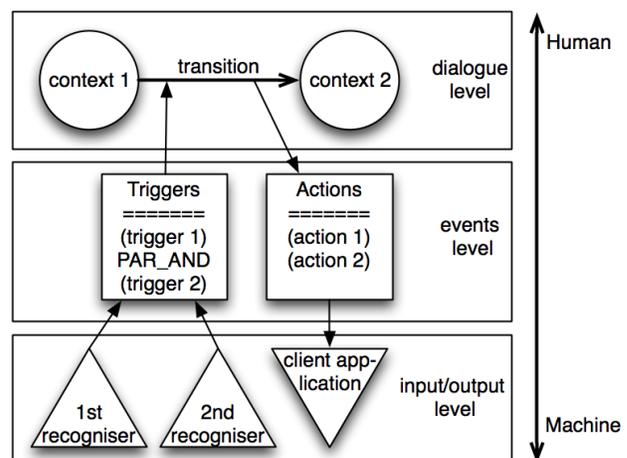


Figure 1. The three levels of SMUIML

The spectrum of multimodal dialogue description language users, on a scale from usability to expressiveness, was presented in [8]. Through various workshops, informal discussions with colleagues and students and a study of the current state of the art, we envisioned three types of approaches for

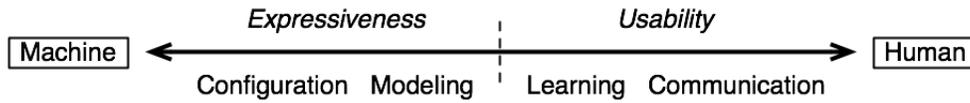


Figure 2. Four purposes of modelling languages.

a description language: a highly formal language approach that perfectly fits for configuring a tool, a loosely formal language approach which is good for communicating the details of an application and a “middle” approach focussing on the modelling. Along these three approaches, a formal language can also be used as a learning tool as shown in Figure 2. As a learning tool, a modelling language helps teachers in communicating the features of a particular application domain to their students.

In [8] we presented 9 guidelines for a multimodal description language. These guidelines should be used as design tools or as language analysis criteria:

- Abstraction levels
- Modelling the human-machine dialogue
- Adaptability to context and user (input and output)
- Control over fusion mechanism
- Control over time synchronicity
- Error handling
- Event management
- Input and output sources representation
- Finding the right balance between usability and expressiveness

SOFTWARE SUPPORT FOR SMUIML

As a language defining a full model of multimodal human-machine events and dialogues, SMUIML already provides modelling capabilities as well as a reflection basis. However, the language shows its true potential when linked to a range of different supportive software solutions. In the following, we briefly introduce software support within SMUIML for interpretation and then discuss the latest software addition in the form of a graphical editor for designing multimodal human-machine dialogues.

The HephaisTK Framework

A framework for the creation of multimodal interfaces has been developed in our research lab. This framework, called HephaisTK, uses SMUIML as its scripting language. Note that it is out of the scope of this paper to provide a detailed description of HephaisTK. However, the interested reader can find more details about HephaisTK in [7].

Basically, a description created in SMUIML with the structure shown in Listing 1 is used to configure the HephaisTK framework. The `<recognizers>` part indicates which

recognisers have to be loaded by the framework. It further provides some high-level parameters such as whether a speech recogniser is able to recognise different languages. The `<triggers>` are directly bound to the different fusion algorithms provided by HephaisTK. The `<actions>` part defines the semantics to be used when communicating fusion results to a client application. Last but not least, the `<dialog>` part of SMUIML is used in HephaisTK for a number of specific goals.

Listing 1. Basic layout of a SMUIML script

```
<?xml version="1.0" encoding="UTF-8"?>
<smuiml>
  <integration_desc client="client_app">
    <recognizers>
      <!-- ... -->
    </recognizers>
    <triggers>
      <!-- ... -->
    </triggers>
    <actions>
      <!-- ... -->
    </actions>
    <dialog>
      <!-- ... -->
    </dialog>
  </integration_desc>
</smuiml>
```

First and foremost, by providing a description of the human-machine dialogue flow, the HephaisTK `DialogManager` agent stays in a consistent state with the client application. The clear separation of the SMUIML `<dialog>` into transitions and contexts allows the different triggers to be enabled or disabled depending of the current context. Since only a subset of triggers has to be considered in a given context, the load on the recognisers is reduced and the overall recognition rate is improved. The `<dialog>` part of SMUIML also helps with the instantiation of the different fusion algorithms present in HephaisTK. In the case of the Hidden Markov Model-based fusion algorithm that is integrated in HephaisTK, the definition of the human-machine dialogue in SMUIML is also used to generate a set of all expected trigger input sequences. The set of all expected sequences is then injected into a series of Hidden Markov Models (one per context of use), so that the fusion engine is ready to be used when launching the HephaisTK framework.

In the context of the HephaisTK framework, SMUIML is applied at multiple levels: at the multimodal dialogue description level, at the recogniser launch and parameterisation level as well as at the fusion engine instantiation level. Thus, SMUIML is typically used during the later stages of multimodal interface development, i.e. at the *system design* and *runtime* stages.

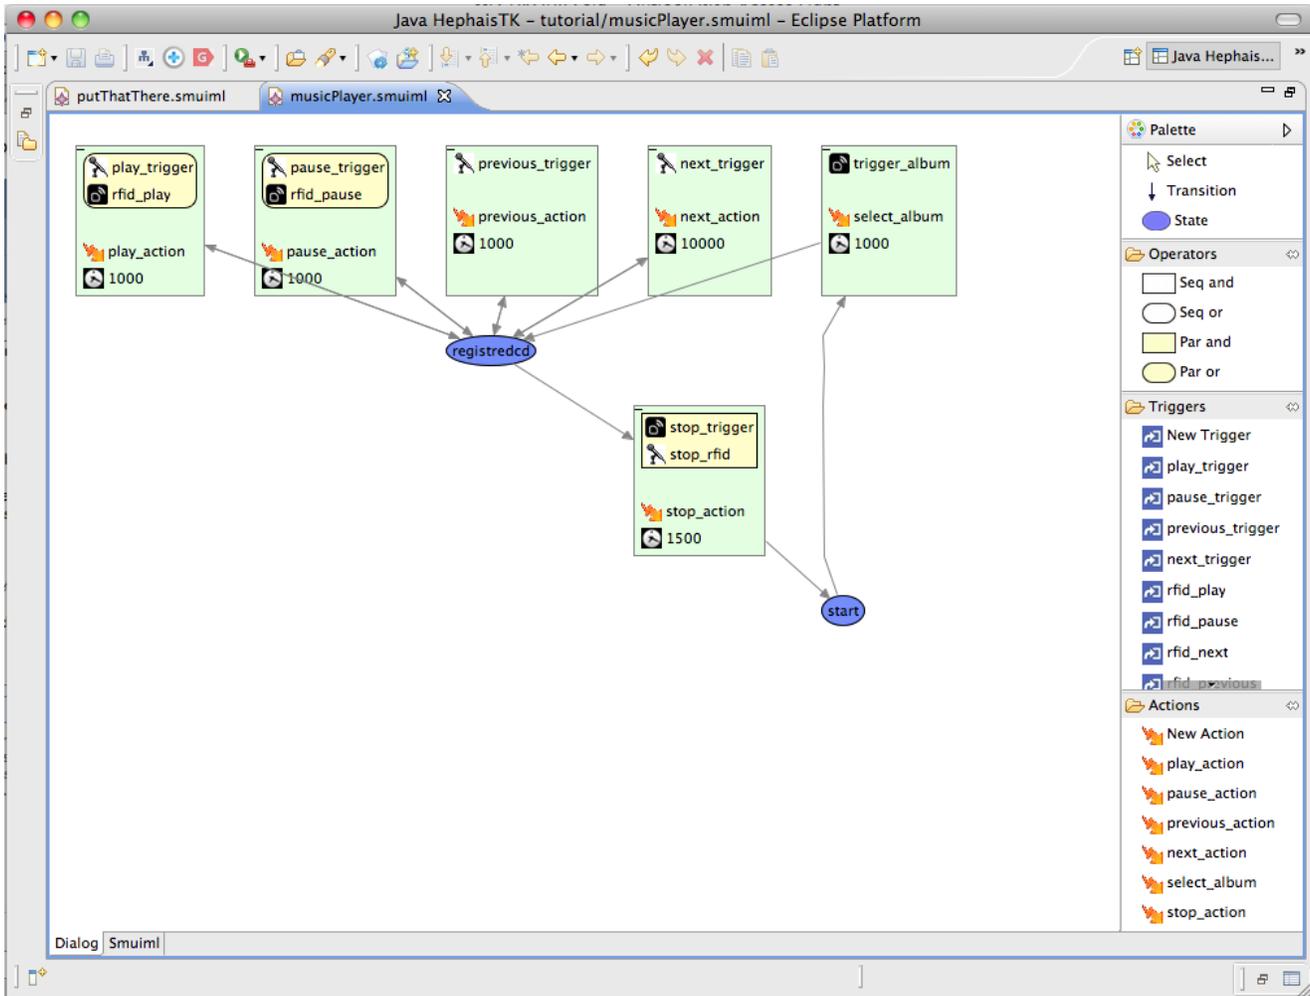


Figure 3. The SMUIML graphical editor

The SMUIML Graphical Editor

As SMUIML is a language that is derived from the XML metalanguage, a standard text editor is sufficient for creating SMUIML documents. Even if the language has been proven to be expressive in a qualitative study [8], the editing of “raw” XML easily leads to errors which can only be identified when interpreting a SMUIML document at runtime. Other issues with the text-based editing of SMUIML scripts include the lack of an explicit representation of the relationships between different elements as well as the difficulty to produce and maintain an accurate mental model of complex dialogue scenarios. Furthermore, the necessity of having to learn a new language can represent a serious challenge for some users. In order to overcome these shortcomings, we have developed a graphical editor for the definition of new SMUIML scripts.

The goal of our SMUIML graphical editor was to provide a usable and expressive tool for creating SMUIML scripts for HephaisTK to developers who are not fully proficient with multimodal interfaces. The dialogue editor offers a graphical representation of SMUIML-encoded multimodal human-

machine dialogues. Furthermore, it supports the creation of sets of actions and triggers and can be used to generate the Java configuration file with all calls related to the SMUIML script. As shown later, the graphical representation of a multimodal dialogue closely follows the SMUIML logic presented in the previous section. The SMUIML graphical editor tool was created by using the Eclipse⁵ open development platform as a basis. Eclipse is widely used and known among development teams and its usage provides expert users with a set of well-known interface elements. The SMUIML graphical tool itself was created using the Graphical Editing Framework (GEF)⁶ and the Eclipse Modeling Framework (EMF)⁷.

The main window of the graphical editor is shown in Figure 3. The central part of the tool is dedicated to the actual dialogue representation. As stated before, the multimodal human-machine dialogue in SMUIML typically takes the form of a state machine. A graphical representation of this state machine is used to depict the multimodal dialogue in

⁵<http://www.eclipse.org>

⁶<http://www.eclipse.org/gef/>

⁷<http://www.eclipse.org/modeling/emf/>

the graphical editor. Note that the editor also provides access to a textual version of the SMUIML script that is currently edited. Any changes done either in the graphical or the textual representation of a given SMUIML script are reflected in the other representation. The checking of errors is achieved in real-time for both the graphical and textual representations.

On the right-hand side of the window are a set of toolboxes and most of them are related to the different parts of a typical SMUIML file. The `Palette` toolbox presents the basic building blocks for creating the dialogue state machine, in particular states and transitions. The selection tool is also presented in the `Palette` toolbox. The `Operators` toolbox offers some operators for the combination of different modalities as defined in the SMUIML specification. Note that these operators are tightly linked to the CARE properties [6]. `Seq` and corresponds to sequential-constrained complementarity, `Par` and to sequential-unconstrained complementarity, `Seq` or to equivalence and `Par` or to redundancy. The next toolbox is devoted to input triggers and contains a list of all triggers defined for a given application, as well as a `New trigger` button to create new triggers. Last but not least, the `Actions` toolbox list all actions that have been defined for a given application and also provides a `New action` button. Triggers and actions are added to these toolboxes as they are defined along the creation of the multimodal dialog.

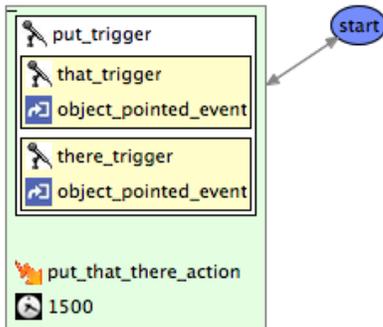


Figure 4. Graphical description of “put that there”

Figure 4 shows the graphical representation of Richard Bolt’s infamous “put that there” example [3] in the graphical editor whereas the corresponding textual SMUIML specification is shown in Listing 2. In the same way as in Figure 3, states (contexts) are visualised as blue ellipses. Along the actions taken by users, HephaisTK might stay in the same context or switch to another context of use. In the “put that there” example, there is only as single `start` context with a transition starting and pointing to it. This transition contains the overall description of the “put that there” action. It asks for five different input triggers in order that the action will be fired. Namely, three speech triggers (“put”, “that” and “there”), as well as two pointing event triggers. Furthermore, three temporal combination operators are used in this example. The main transition uses a `Seq` and operator asking for a “put” speech trigger to be followed by a “that” and “there” sub-event. Those two sub-events use a `Par` and combina-

tion operator which means that there should be speech and pointing triggers but without any sequential constraint. This implies that a user can perform the commands in different orders, such as “put that” [point1] [point2] “there” or “put” [point1] “that there” [point2]. Both sequences will be correctly recognised by HephaisTK. Finally, the transition specifies a time window of 1500 milliseconds for the whole command as well as an action, in the form of a message to be sent the client application, to be performed if the command has been successfully recognised. The transition in our example then proceeds to the same `start` context it originated from but of course we could also move to any other defined context.

The SMUIML graphical editor has been presented to two expert users in order to achieve a small expert review. This review by experts lead to a number of changes to improve the editor’s usability. The modality of each trigger is now indicated by means of an icon. The `start` context which is the only mandatory context in a given SMUIML script is visualised in a slightly different colour to denote its special status compared to the other contexts. Finally, users have the possibility to change the colour of connections, contexts or transitions in order to best suit their preferences.

Listing 2. SMUIML description of the “put that there” example

```
<context name="start">
<transition leadtime="1500">
  <seq_and>
    <trigger name="put_trigger"/>
    <transition>
      <par_and>
        <trigger name="that_trigger"/>
        <trigger name="object_pointed_event"/>
      </par_and>
    </transition>
    <transition>
      <par_and>
        <trigger name="there_trigger"/>
        <trigger name="object_pointed_event"/>
      </par_and>
    </transition>
  </seq_and>
  <result action="put_that_there_action"/>
  <result context="start">
</transition>
</context>
```

The graphical editor also contains an integrated debugging tool. This debugging tool is launched along the client application and indicates in real-time the context the HephaisTK framework is currently in. It also highlights the transition that lead from the previous context to the current one. In the example illustrated in Figure 5, the application starts in the `start` context. An RFID reader that is connected to the framework detects a tagged music album and transmits the information. Based on the `trigger_album` trigger, a transition is fired and the application moves to the `registeredcd` state and starts playing the music album. Then the user utters a `stop` command and holds at the same time an RFID tag labelled “stop” close to the RFID reader. This simultaneous action fires the transition going from the `registeredcd` context back to the `start` context. As illustrated in this example, the graphical debugging tool al-

allows developers to visually analyse the behaviour of an application in real-time.

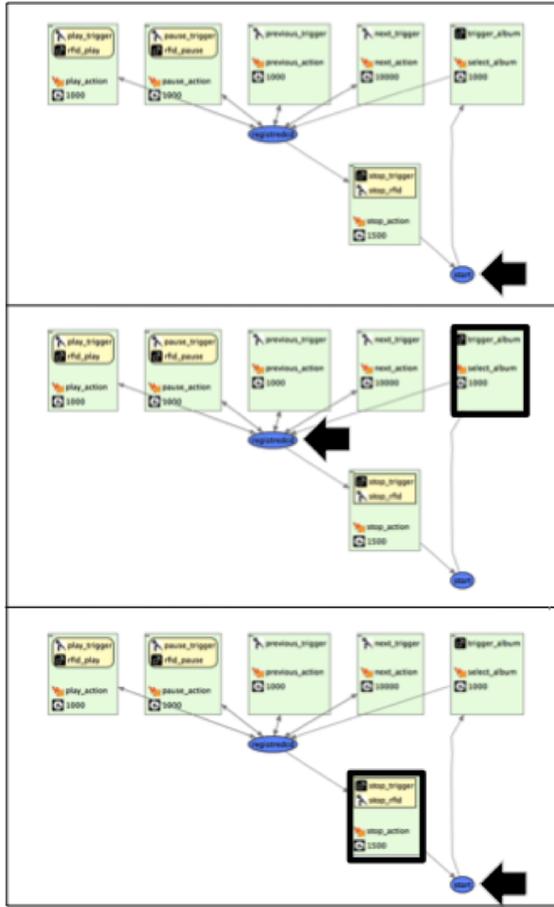


Figure 5. Graphical debugging tool

FUTURE WORK

While the presented SMUIML graphical editor looks promising and offers some features not available in other graphical editors for creation of multimodal interfaces, we plan to perform a detailed evaluation of the presented solution in the near future. First, we are going to evaluate the usability of the presented graphical editor for multimodal interaction design. However, in addition to the usability of the graphical editor, we also plan to evaluate the expressiveness of the presented approach. It is not enough to guarantee an effective and simple definition of multimodal interactions based on the graphical editor. We also have to ensure that the editor and the underlying SMUIML language are expressive enough to describe multimodal interactions of arbitrary complexity.

Another important future direction is to support the flexible adaptation of multimodal interfaces. The idea is to no longer have a fixed combination of modalities, but rather provide a context-dependent adaptation of multimodal user interfaces. This can either be achieved by extending the SMUIML language with the necessary concepts or by introducing another language for the adaptation of the multimodal interaction. In

this view, the abstract user interface definition would rely on SMUIML while the concrete, context-dependant user interface specification would be based on a new language. The final user interface would be realised by HephaisTK [5].

Last but not least, multimodal interfaces should not be treated isolated from the underlying content that they are providing access to. Today's document formats do often not offer access to specific semantic subparts and media types embedded in a given document [14]. If we are able to model and get access to information at a finer granularity level, the different modalities of interaction can also be associated with specific parts of a document. Within the MobiCraNT project, we are currently investigating innovative mobile cross-media applications. As part of this research effort, we are developing a new fluid cross-media document model that offers the fine grained access to information described earlier. By no longer defining the modality of interaction at the document but rather at the level of cross-media information units, we aim for a more flexible solution where the structure of a document as well as the multimodal interface to interact with its parts is adapted in a content-dependent manner.

CONCLUSION

We have presented our exploration on software support for multimodal UIDL based on the SMUIML multimodal dialogue modelling language. Thereby, we focussed on two particular software components: the HephaisTK framework which is used to interpret the SMUIML language (at the *runtime* stage) and the SMUIML graphical editor for the graphical design of multimodal interaction dialogues (at the *system design* stage). The SMUIML graphical editor aims to provide a user-friendly way to create multimodal applications based on HephaisTK and SMUIML. Compared to other graphical dialogue editors, our solution supports temporal constraints and a number of operators for the combination of multiple modalities. While these concepts already form part of the underlying SMUIML language, the graphical editor makes these concepts accessible via a user-friendly interface. Users further have the possibility to freely switch between the graphical and textual dialogue representation. The presented graphical editor for SMUIML further introduces a number of usability-oriented features such as automatic layouting, the clear identification of input modalities via specific icons as well as the possibility to customise various features of the graphical editor. Last but not least, the SMUIML graphical editor offers an integrated debugging tool supporting developers in analysing the real-time behaviour of their applications.

ACKNOWLEDGMENTS

The authors wish to thank Saïd Mechour for his work on the SMUIML graphical editor. The work on HephaisTK and SMUIML were funded by the Hasler Foundation in the context of the MeModules project and by the Swiss National Center of Competence in Research on Interactive Multimodal Information Management via the NCCR IM2 project. Bruno Dumas is supported by the MobiCraNT project in the context of the Strategic Platforms programme from the Brussels Institute for Research and Innovation (Innoviris).

REFERENCES

1. M. Araki and K. Tachibana. Multimodal Dialog Description Language for Rapid System Development. In *Proceedings of the 7th SIGdial Workshop on Discourse and Dialogue*, pages 109–116, Sydney, Australia, July 2006.
2. J. D. Boeck, D. Vanacken, C. Raymaekers, and K. Coninx. High-Level Modeling of Multimodal Interaction Techniques Using NiMMiT. *Journal of Virtual Reality and Broadcasting*, 4(2), September 2007.
3. R. A. Bolt. “Put-that-there”: Voice and Gesture at the Graphics Interface. In *Proceedings of the 7th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '80)*, pages 262–270, Seattle, USA, July 1980.
4. M.-L. Bourguet. A Toolkit for Creating and Testing Multimodal Interface Designs. In *Adjunct Proceedings of the 15th Annual Symposium on User Interface Software and Technology (UIST 2002)*, Paris, France, October 2002.
5. G. Calvary, J. Coutaz, D. Thevenin, Q. Limbourg, L. Bouillon, and J. Vanderdonckt. A Unifying Reference Framework for Multi-Target User Interfaces. *Interacting with Computers*, 15(3):289–308, 2003. Computer-Aided Design of User Interface.
6. J. Coutaz, L. Nigay, D. Salber, A. Blandford, J. May, and R. M. Young. Four Easy Pieces for Assessing the Usability of Multimodal Interaction: The CARE Properties. In *Proceedings of the 5th International Conference on Human-Computer Interaction (Interact 1995)*, Lillehammer, Norway, June 1995.
7. B. Dumas, D. Lalanne, and R. Ingold. HephaisTK: A Toolkit for Rapid Prototyping of Multimodal Interfaces. In *Proceedings of 11th International Conference on Multimodal Interfaces (ICMI 2009)*, pages 231–232, Cambridge, USA, November 2009.
8. B. Dumas, D. Lalanne, and R. Ingold. Description Languages for Multimodal Interaction: A Set of Guidelines and its Illustration with SMUIML. *Journal on Multimodal User Interfaces: “Special Issue on The Challenges of Engineering Multimodal Interaction”*, 3(3):237–247, February 2010.
9. K. Katsurada, Y. Nakamura, H. Yamada, and T. Nitta. XISL: A Language for Describing Multimodal Interaction Scenarios. In *Proceedings of the 5th International Conference on Multimodal Interfaces (ICMI 2003)*, pages 281–284, Vancouver, Canada, November 2003.
10. W. A. König, R. Rädle, and H. Reiterer. Squidy: A Zoomable Design Environment for Natural User Interfaces. In *Proceedings of the 27th International Conference on Human Factors in Computing Systems (CHI 2009)*, pages 4561–4566, Boston, USA, April 2009.
11. J.-F. Ladry, P. Palanque, S. Basnyat, E. Barboni, and D. Navarre. Dealing with Reliability and Evolvability in Description Techniques for Next Generation User Interfaces. In *Proceedings of the 26th ACM International Conference on Human Factors in Computing Systems (CHI 2008)*, Florence, Italy, April 2008.
12. J.-Y. L. Lawson, A.-A. Al-Akkad, J. Vanderdonckt, and B. Macq. An Open Source Workbench for Prototyping Multimodal Interactions Based on Off-the-Shelf Heterogeneous Components. In *Proceedings of the 1st ACM SIGCHI Symposium on Engineering Interactive Computing Systems (EICS 2009)*, pages 245–254, Pittsburgh, USA, July 2009.
13. F. Paternò, C. Santoro, J. Mäntyjärvi, G. Mori, and S. Sansone. Authoring Pervasive Multimodal User Interfaces. *International Journal of Web Engineering and Technology*, 4(2):235–261, 2008.
14. B. Signer. What Is Wrong with Digital Documents? A Conceptual Model for Structural Cross-Media Content Composition and Reuse. In *Proceedings of the 29th International Conference on Conceptual Modeling (ER 2010)*, pages 391–404, Vancouver, Canada, November 2010.
15. A. Stanculescu, Q. Limbourg, J. Vanderdonckt, B. Michotte, and F. Montero. A Transformational Approach for Multimodal Web User Interfaces Based on UsiXML. In *Proceedings of the 7th International Conference on Multimodal Interfaces (ICMI 2005)*, pages 259–266, Toronto, Italy, October 2005.