Journal of Universal Computer Science, vol. 8, no. 10 (2002), 871-880
submitted: 24/4/02, accepted: 15/10/02, appeared: 28/10/02 © J.UCS

Research in Structural Computing

David L. Hicks
(Department of Computer Science
Aalborg University Esbjerg
hicks@cs.aue.auc.dk)

Uffe K. Wiil
(Department of Computer Science
Aalborg University Esbjerg
ukwiil@cs.aue.auc.dk)

Peter J. Nürnberg
(Department of Computer Science
Aalborg University Esbjerg
pnuern@cs.aue.auc.dk)

Abstract: Structural computing is one of the most recent research threads to emerge in the field of hypermedia. Though a relatively new line of study, research results have already started to emerge in the structural computing field. This paper examines a number of structural computing research projects to provide an overview of the current state of the field as well as a look at the direction of ongoing projects. It also briefly discusses additional areas of research in structural computing that will be important to consider as research in the field continues.

Keywords: Structural Computing, Open Hypermedia, Hypermedia Architectures
Categories: H.5.4, D.2.11

1 Introduction

Structural computing is one of the latest research threads to emerge in the field of hypermedia. Though many of the ideas inherent in structural computing originated in the open hypermedia systems area of research, structural computing is actually a confluence of a number of research threads. Its component research areas include the identification and investigation of alternative hypermedia domains, the study of structuring mechanisms and the important roles they play within application domains, and the design and development of hypermedia architectures.

This paper examines research in the structural computing field. It provides an overview of existing research results that have been reported in the structural computing area along with a look at the direction of ongoing research projects. The paper continues in the following section with background material that briefly defines structural computing and describes the developments that lead to the establishment of the field. The third section examines several specific structural computing projects to provide an overview of the existing research along with an indication of the direction in which ongoing research is leading. The discussion in the fourth section characterizes the structural computing research results that have been reported to date, along with a look at important areas to consider as structural computing research continues. A brief summary then concludes the paper.

2 Background

Many of the ideas that lead to the formulation of the structural computing field originated in the open hypermedia systems area of research. Research into open hypermedia systems has been ongoing for over a decade. Motivated by the problems identified with early hypermedia systems [Halasz 1988], open hypermedia researchers proposed new approaches to providing hypermedia functionality. Openness was a key characteristic of these new hypermedia architectures. An open hypermedia system is one that makes the services it provides available through a well-defined interface and offers them in a flexible and modular way. These characteristics enable the hypermedia services offered to be used by an open set of client applications. Substantial progress has been made over the years in the open hypermedia research area. Successful example systems include DHM [Grønbæk, et al. 1994], Microcosm [Hall, et al. 1996], Chimera [Anderson, et al. 1994], and many others.

Despite the success of open hypermedia systems, as work progressed in the area, researchers began to notice a limiting characteristic that these systems all shared. The functionality the systems provided was specifically designed to support a particular type of hypermedia - referential (or navigational) hypermedia. Referential hypermedia, the ability to create associative links between information items, was certainly a useful capability, as demonstrated by the success of many existing hypermedia systems. However, new types of hypermedia functionality (or hypermedia domains) were being identified such as spatial hypermedia [Marshall and Shipman 1994] and taxonomic hypermedia [Parunak 1991]. These new hypermedia domains represented important capabilities that open hypermedia researchers wanted to provide for the users of their systems. However, the operations needed to support these new hypermedia domains could not be easily or naturally mapped to the navigational hypermedia model upon which open hypermedia systems were based. As the set of hypermedia domains continued to expand, the referential bias of these systems became increasingly apparent and limiting [Anderson 2000].

Researchers began to realize that in order to effectively support multiple hypermedia domains, they required a new model upon which to base their systems, a model that was capable of supporting a range of hypermedia (or structure) domains in addition to the referential one. As they searched for an appropriate model, researchers noticed that the prominence of structure was the single common characteristic shared across all hypermedia domains. For example, structure was critical to defining associative connections in referential hypermedia applications. Similarly, structure was what enabled taxonomic hypermedia to be used to bring order to a collection of objects, or for spatial hypermedia techniques to be used to organize an emerging idea space. This observation had a significant influence, and lead to proposals for a new type of hypermedia paradigm, one in which structure played a much more prominent role. The term structural computing is used to describe environments based on this new paradigm, reflecting the importance of structure [Nürnberg et al. 1997].

Figure 1 illustrates conceptually a generic structural computing environment. Although similar in appearance to an open hypermedia environment, significant differences distinguish structural computing systems. Unlike predecessor systems, structure plays a much more fundamental role in a structural computing environment. It serves as the fundamental building block upon which all abstractions in the environment are based. This is illustrated in Figure 1 by the presence of a structure store at the lowest level of the architecture. It provides a basic unit of structure, the structural atom, and all other abstractions are built upon it. In the middle layer of Figure 1 are structure servers. A structure server is an entity that builds upon the basic structure store services to implement a specific type of abstraction, and to make that abstraction available to client applications. For example, a structure server could be designed to provide the abstractions necessary to support a particular hypermedia domain, such as the referential or spatial hypermedia domain. The functionality of structure servers is not intended for use directly by end users. It is instead integrated into applications that build upon the functionality to offer useful capabilities to end users.

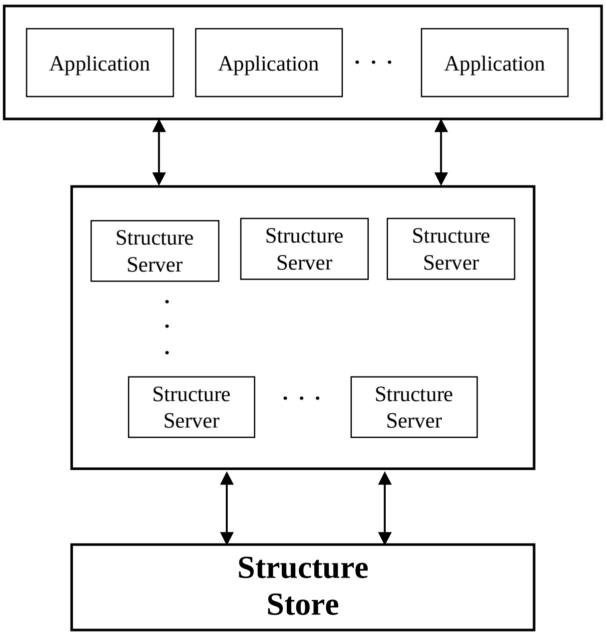

Figure 1: A Generic Structural Computing Environment.

3 Structural Computing Research Projects

Though a relatively new area of study, research results have already begun to emerge for the structural computing area [Nürnberg 1999, Reich and Anderson 2000, Reich et al. 2001]. Research has started on a variety of fronts including the design and development of structural computing environments, the investigation of structure domain interoperability, and the use of structural computing techniques and technology to support specific application areas. This section will examine a number of research projects in the structural computing area. The intent is to provide a representative sample of the research results that have been reported to date as well as the projects that are currently underway in the structural computing area.

- The HOSS project was conducted by researchers at Texas A&M University. It proposed a view of hypermedia that considered it to be a new computing paradigm [Nürnberg et al. 1996], thereby laying the groundwork for much of what was to become structural computing. The HOSS project sought to integrate hypermedia (structure) into the lowest, most fundamental levels of a computing environment - the operating system. It argued for an augmentation of standard operating system services to include functionality specifically intended to support structural entities. This enabled structural objects to exist at the operating system level, and to be treated in a way that was analogous to data objects. Policies were also developed that enabled operating systems to benefit from an awareness of structure. For example, caching and access control policies of the operating environment were augmented to handle structural objects as well as data objects in an appropriate way. The ability of the operating system to understand and process structure enabled the development of facilities such as the prefetching of information (hypermedia nodes) based on semantic locality (structural connectedness) [Nürnberg et al. 1996].

- Researchers at the University of Patras have proposed a computing environment, based on a component-based open hypermedia system, that can support a range of useful structural abstractions [Vaitis et al. 2000]. The objective is to support not only a variety of hypermedia application domains such as the navigational, spatial, and taxonomic domains, but also provide support for specialization within each domain (sub-domains) and the combination of domains (applications). Using a template mechanism, structural patterns can be defined that establish a structural model upon which structure services can operate. The generic structural building block is the Abstract Structural Element. Templates enable specializations of the Abstract Structural Element to be defined to create the specialized group of structural elements that are required to support a particular structural domain. This approach results in an intensely structure oriented middleware layer of the component based open hypermedia system. In addition, the structured modeling of domains (and sub-domains or applications) inherent in the approach encourages a thorough analysis and understanding of the domain, and also promotes the reuse of existing models when developing support for a new domain [Kyriakopoulou et al. 2001].

- The Construct system is a component-based open hypermedia system that builds upon previous research performed at Texas A&M University, Aalborg Unversity, and Aarhus University. The researchers collaborating on the Construct project have experience with the development of several previous open hypermedia systems, most notably HOSS [Nürnberg et al. 1996], DHM [Grønbæk and Trigg 1994], and HyperDisco [Wiil and Leggett 1996]. A primary goal of the project is to combine the best aspects of these previous systems as well as other existing ones to create a comprehensive environment for supporting structure. An important aspect of Construct is that it has from the start been specifically designed to be compliant with the standards proposed by the Open Hypermedia Systems Working Group (OHSWG) [Davis, et al. 1999, Reich et al. 1999]. An implementation of the Construct system has been underway at Aalborg University Esbjerg since 1998. Currently it provides services to support several hypermedia domains including the navigational, taxonomic, and spatial domains. In addition, facilities are provided to support metadata creation and management, data mining, and cooperation [Wiil and Hicks 2001]. Tools are also provided in the Construct environment to assist in the development of structure services. Currently these include the UML Tool, which allows new structure services to be specified in the UML language, and the CSC (Construct Service Compiler), which can process structure service specifications and generate code for them [Wiil 2000, Wiil et al. 2000].

- The Fundamental Open Hypermedia Model (FOHM) was developed at the University of Southampton [Millard et al. 2000]. FOHM is based on, and is an extension of the OHP (open hypermedia protocol), a standard that was developed by the open hypermedia systems working group to support the interoperability of hypermedia systems. OHP was specifically targeted to support interoperability within the navigational hypermedia domain. FOHM broadens its applicability to also include the spatial, and taxonomic hypermedia domains [Millard and Davis 2000]. FOHM defines a data model and set of related operations that are capable of representing the structural abstractions and operations of multiple hypermedia domains. Representing each of the domains within a single model provides a way in which the structures of one domain can be understood within the context of another. This enables and promotes interoperability between domains. For example, in a multi domain hypermedia environment, a user browsing navigational hypermedia structures might encounter a spatial hypermedia structure. In the FOHM environment, a "spatial structure aware" navigational browser might be able to interpret spatial hypermedia structures and enable them to be browsed as if they were navigational structures. In addition to this basic notion of domain interoperability, a cross fertilization effect between domains is possible. Specifically, the model can support transformations that enable useful features within one domain to be used within another domain in which the feature is not inherently present. For example, the notion of anchors as known in the navigational hypermedia domain provide the ability to link together arbitrarily sized sections or elements of information items (documents). This granularity scaling capability could be extended to the spatial domain to enable spatial hypermedia organizational functionality to be applied to arbitrary (finer) granularities, e.g., to allow the parts or sub-parts of larger documents to be organized spatially [Millard et al. 2000].

- At the University of Colorado researchers have focused on the application of structural computing technologies and techniques to a particular application area, software engineering. They have analyzed a number of software engineering subdomains in order to generate a general set of requirements that structural computing environments must meet in order to accommodate the software engineering area [Anderson 1999]. Subsequently, a more in depth look has been taken at how structural computing can support the specific software engineering subdomain of information integration. Structural computing has influenced the development of the InfiniTe Information Integration Environment [Anderson and Sherba 2001]. In particular, the design of the InfiniTe information repository facility is based upon structural computing principles. It is being implemented as a set of structure servers that provide a specific set of structural abstractions in order to create a scalable and flexible solution which can be adapted as needed to include new information integration capabilities.

- Research at the Fraunhofer IPSI Institute (formerly GMD-IPSI) has been conducted that focuses on the integration of the workflow hypermedia domain into structural computing environments [Wang and Haake 1999]. The primary focus has been the development of a structure service component that can support the workflow process within applications. Deployment of the component in a structural computing environment enables it to utilize functionality from other domains, e.g., navigational hypermedia. This facilitates the integration of heterogeneous information systems as well as heterogeneous workflow systems. A related project has examined the use of a structural computing environment to support collaboration [Haake 2000]. The specific focus has been on supporting the collaboration subtask of coordination by using structures to model the entities that comprise a shared workspace: processes, teams, and content. The hope is that by modeling a workspace with structures it will be possible to define coordination support as computations over the structures, and that this will provide a more flexible coordination mechanism that can be adapted to the changing requirements inherent in cooperative work processes. An additional project at the Fraunhofer IPSI Institute involves user interface issues of structural computing [Wang and Fernandez 2001]. As part of the project, a graphical hypermedia user interface has been designed that integrates features from the navigational, spatial, taxonomic, and workflow domains.

4 Discussion

The projects summarized in the previous section demonstrate the diversity of research in the structural computing area. Though a variety of issues are being examined, current structural computing research projects can be roughly categorized into one of three areas according to their primary focus. Many of the projects involve alternative hypermedia domain related issues. Examples include models for supporting multiple hypermedia domains [Millard et al. 2000] and the development of multiple domain graphical hypermedia interfaces [Wang and Fernandez 2001]. Other projects have centered on the development of environments for supporting structural computing [Vaitis et al. 2000, Wiil and Hicks 2001]. Most of these have focused on how existing hypermedia architectures, especially open hypermedia architectures, can be transformed into environments for supporting structural computing. A third area being considered is the use of structural computing environments to support specific application areas. Examples include using structural computing techniques and technology to support software engineering [Anderson 1999, Anderson and Sherba 2001], workflow [Wang and Haake 1999], collaboration [Haake 2000], literary systems [Rosenberg 2000], and scholarly discourse [Buckingham-Shum 2000].

The issues being investigated in each of these areas of structural computing research are important ones, as indicated by the research attention they have received. Recently, suggestions have been made that an additional area of research be pursued within the structural computing field [Hicks 2000]. Specifically, some have suggested the need to focus research attention on the lower levels of structural computing environments (Figure 1). For example, the storage component at the lowest level of a structural computing environment is a *structure store*. This is one of the defining characteristics of a structural computing environment. However, in most structural computing prototype environments, the storage component residing at the lowest level closely resembles the storage component of open hypermedia systems - a data oriented storage component rather than a structure oriented one. Few current research project prototypes are actually based on a structure store at the lowest level [Vaitis et al. 2000, Wiil and Hicks 2001].

There are many issues that need to be investigated related to the lower levels of structural computing environments. Many of these issues pertain to the structure store itself. For example, what operations should the structure store provide? How should the functionality it provides be offered to applications - what should the interface look like? What is the appropriate model of structure upon which to build a structure server? By focusing research attention on these and related lower level issues, progress and improvements can be expected in the structuring capabilities these systems provide. This, in turn, will help to advance research at all levels within structural computing environments.

5 Summary

Structural computing is one of the newer areas of research within the field of hypermedia. It was formulated as the confluence of several existing research threads including the investigation of alternative hypermedia domains, the study of structuring mechanisms within applications, and the design of hypermedia architectures. A number of structural computing research projects are underway investigating various aspects of structural computing. Most have focused on one of three primary areas: issues concerning alternative hypermedia domains, the development of environments for supporting structural computing, and the use of structural computing technology to support specific application areas. These are important areas to consider, as demonstrated by the research attention they have received. However, as the structural computing field moves forward, it is important to consider additional threads of research as well. Important among these new areas to consider are issues concerning the lower levels of structural computing environments. Focusing attention on such fundamental issues holds significant potential to benefit the field as a whole.

References

[Anderson, et al. 1994] Anderson, K., Taylor, R., and Whitehead, E. J. 1994. Chimera: Hypertext for heterogeneous software environments. In *Proceedings of the 1994 ACM European Conference on Hypertext*, (Edinburgh, Scotland, Sep), ACM Press, pp. 94-107.

[Anderson 1999] Anderson, K. 1999. Software engineering requirements for structural computing. In Proceedings of the First Workshop on Structural Computing (Darmstadt, Germany, Feb), pp 22-26. Available in technical report AUE-CS-99-04 (Aalborg University Esbjerg, Denmark).

[Anderson and Sherba 2001] Anderson, K. and Sherba, S. 2001. Using Structural Computing to Support Information Integration. In Proceedings of the Third International Workshop on Structural Computing (Århus, Denmark, August) pp. 151-159.

[Anderson 2000] Anderson K. 2000. Introduction to SC2. In Proceedings of the Second International Workshop on Structural Computing (San Antonio, TX, USA, May). Berlin: Springer-Verlag LNCS 1903 pp. 96-97.

[Buckingham Shum 2000] Buckingham Shum, Simon. 2000. Scholarly Discourse as Computable Structure. In Proceedings of the Second International Workshop on Structural Computing (San Antonio, TX, USA, May). Berlin: Springer-Verlag LNCS 1903 pp. 120-128.

[Davis, et al. 1999] H. C. Davis, D. E. Millard, S. Reich, N. Bouvin, K. Grønbæk, P. J. Nürnberg, L. Sloth, U. K. Wiil, and K. Anderson. 1999. Interoperability between Hypermedia Systems: The Standardisation Work of the OHSWG (technical briefing). In Proceedings of the '99 ACM Conference on Hypertext, February 21-25, 1999, Darmstadt, Germany, pages 201-202, February 1999.

[Grønbæk and Trigg 1994] Grønbæk, K. and Trigg, R. H. 1994. Design Issues for a Dexter-based Hypermedia System. Communications of the ACM, 37, 2, (Feb), 40-49.

[Grønbæk, et al. 1994] Grønbæk, K., Hem, J. A., Madsen, O. L., and Sloth, L. 1994. Cooperative hypermedia systems: A Dexter-based architecture. *Communications of the ACM*, 37, 2, (Feb), 64-74.

[Haake 2000] Haake, J. 2000. Structural Computing in the Collaborative Work Domain? In Proceedings of the Second International Workshop on Structural Computing (San Antonio, TX, USA, May). Berlin: Springer-Verlag LNCS 1903, pp. 108-119.

[Halasz 1988] Halasz, F. 1988. Reflections on Notecards: Seven issues for the next generation of hypermedia systems. Communications of the ACM, 31, 7, (July), 836-852.

[Hall, et al. 1996] Hall, W., Davis, H., and Hutchings, G. 1996. *Rethinking Hypermedia – The Microcosm Approach*. Kluwer Academic Publishers.

[Hicks 2000] Hicks, D. L. 2000. Structural Compupting: Evolutionary or Revolutinoary? In Proceedings of the Second International Workshop on Structural Computing (San Antonio, TX, USA, May). Berlin: Springer-Verlag LNCS 1903, pp. 170-178.

[Kyriakopoulou et al. 2001] Kyriakopoulou, M., Avramidis, D., Vaitis, M., Tzagarakis, M., and Christodoulakis, D. 2001. Broadening Structural Computing Systems Towards

Hypermedia Development. In Proceedings of the Third International Workshop on Structural Computing (Århus, Denmark, August), pp. 131-140.

[Marshall and Shipman 1994] Marshall, C. Shipman, F. M, and Coombs, J.H. 1994. VIKI: Spatial Hypertext Supporting Emergent Structure. In Proceedings of the 1994 ACM European Hypertext Conference, (Edinburgh, Scotland, Sep), ACM Press, pp. 13-23.

[Millard and Davis 2000] Millard, D. and Davis, H. 2000. Navigating Spaces: The Semantics of Cross Domain Interoperability. In Proceedings of the Second International Workshop on Structural Computing (San Antonio, TX, USA, May). Berlin: Springer-Verlag LNCS 1903, pp. 129-139.

[Millard et al. 2000] Millard, D., Moreau, L., Davis, H., and Reich, S. 2000. FOHM: A Fundamental Open Hypertext Model for Investigating Interoperability Between Hypertext Domains. In *Proceedings of the 2000 ACM Hypertext Conference*, (San Antonio, TX, Jun.), 266-267. ACM Press.

[Nürnberg et al. 1996] Nürnberg, P. J., Leggett, J. J., Schneider, E. R., and Schnase, J. L. 1996. Hypermedia Operating Systems: A New Paradigm for Computing. In ACM Hypertext '96 Proceedings, (Washington, CD, US, Mar), ACM Press, pp. 194-202.

[Nürnberg et al. 1997] Nürnberg, P. J., Leggett, J. J., and Schneider, E. R. 1997. As We Should Have Thought. In ACM Hypertext ’97 Proceedings, (Southampton, UK, Apr), ACM Press, pp. 96 - 101.

[Nürnberg 1999] Nürnberg, P. J. (ed.). 1999. *Proceedings of the 1st Workshop on Structural Computing.* Technical Report AUE-CS-99-04, Aalborg University Esbjerg.

[Parunak 1991] Parunak, H. 1991. Don't link me in: Set based hypermedia for taxonomic reasoning. In Proceedings of the 1991 ACM Hypertext Conference, (San Antonio, TX, USA, Dec), ACM Press, 233-242.

[Reich et al. 1999] Siegfried Reich, Uffe K. Wiil, Peter J. Nürnberg, Hugh C. Davis, Kaj Grønbæk, Kenneth M. Anderson, David E. Millard, and Joerg M. Haake. 1999. Addressing interoperability in open hypermedia: The design of the open hypermedia protocol. New Review of Hypermedia and Multimedia, 5:207-248, 1999.

[Reich and Anderson 2000] Reich, S. and Anderson, K. (eds.). 2000. Open Hypermedia Systems and Structural Computing (Proceedings of the Sixth International Workshop on Open Hypermedia Systems and the Second International Workshop on Structural Computing). Berlin: Springer-Verlag LNCS 1903.

[Rosenberg 2000] Rosenberg, J. 2000. Domain Requirements for a Cybertext Authoring System. In Proceedings of the Second International Workshop on Structural Computing (San Antonio, TX, USA, May). Berlin: Springer-Verlag LNCS 1903, pp. 98-107.

[Reich et al. 2001] Reich, S., Tzagarakis, M., and De Bra, P. (eds.). 2001. Proceedings of the Third International Workshop on Structural Computing. Lecture Notes in Computer Science. Springer Verlag, Berlin, August 2001.

[Vaitis et al. 2000] Vaitis, M., Papadopoulos, A., Tzagarakis, M., and Christodoulakis, D. 2000. Towards Structure Specification for Open Hypermedia Systems. In Proceedings of the Second International Workshop on Structural Computing (San Antonio, TX, USA, May). Berlin: Springer-Verlag LNCS 1903, pp. 160-169.

[Wang and Haake 1999] Wang, W. and Haake, J. 1999. Supporting Workflow Using the Open Hypermedia Approach. In *Proceedings of the 1st Workshop on Structural Computing* held in conjunction with Hypertext 99, Darmstadt, Germany, pp. 12-17.

[Wang and Fernandez 2001] Wang, W. and Fernandez, A. 2001. A graphical user interface integrating features from different hypertext domains. In Proceedings of the Third International Workshop on Structural Computing (Aarhus, Denmark, August), pp. 141-150.

[Wiil and Leggett 1996] Wiil, U. K. and Leggett, J. J. 1996. The HyperDisco Approach to Open Hypermedia Systems. In ACM Hypertext '96 Proceeding, (Washington, CD, US, Mar), ACM Press, pp. 140-148.

[Wiil 2000] Wiil, U. K. 2000. Using the Construct Development Environment to Generate a File-Based Hypermedia Storage Service. In Proceedings of the Second International Workshop on Structural Computing (San Antonio, TX, USA, May). Berlin: Springer-Verlag LNCS 1903, pp. 147-159.

[Wiil et al. 2000] Wiil, U. K., Nürnberg, P. J., Hicks, D. L., and Reich, S. 2000. A development environment for building component-based open hypermedia systems. In *Proceedings of the 2000 ACM Hypertext Conference*, (San Antonio, TX, Jun.), 266-267. ACM Press.

[Wiil and Hicks 2001] Wiil, U. K., and Hicks, D. L. 2001. Tools and services for knowledge discovery, management and structuring in digital libraries. In *Proceedings of the Eighth ISPE International Conference on Concurrent Engineering: Research and Applications*, (Anaheim, CA, July), pp. 580-589.