

Howard Williams
Lachlan MacKinnon (Eds.)

LNCS 3112

Key Technologies for Data Management

21st British National Conference on Databases, BNCOD 21
Edinburgh, UK, July 2004
Proceedings




Springer

Howard Williams Lachlan MacKinnon (Eds.)

Key Technologies for Data Management

21st British National Conference on Databases, BNCOD 21
Edinburgh, UK, July 7-9, 2004
Proceedings



Volume Editors

Howard Williams
Lachlan MacKinnon
Heriot-Watt University
School of Mathematical and Computer Sciences
Riccarton, Edinburgh EH14 4AS, UK
E-mail: {mhw, lachlan}@macs.hw.ac.uk

Library of Congress Control Number: 2004108031

CR Subject Classification (1998): H.2, H.3, H.4

ISSN 0302-9743

ISBN 3-540-22382-7 Springer-Verlag Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer-Verlag. Violations are liable to prosecution under the German Copyright Law.

Springer-Verlag is a part of Springer Science+Business Media

springeronline.com

© Springer-Verlag Berlin Heidelberg 2004

Printed in Germany

Typesetting: Camera-ready by author, data conversion by PTP-Berlin, Protago-TeX-Production GmbH
Printed on acid-free paper SPIN: 11017479 06/3142 5 4 3 2 1 0

Preface

This year marked the coming of age of the British National Conference on Databases with its 21st conference held at Heriot-Watt University, Edinburgh, in July 2004. To mark the occasion the general theme of the conference was “When Data Is Key”, reflecting not only the traditional key awarded on a 21st birthday, but also the ever-growing importance of electronic data management in every aspect of our modern lives. The conference was run as part of DAMMS (Data Analysis, Manipulation, Management and Storage) Week, which included a number of co-located and complementary conferences and workshops, including the 2nd Workshop on Teaching, Learning and Assessment in Databases (TLAD2), the BNCOD BioInformatics Workshop, and the 1st International Conference on the Future of Consumer Insight Developments in Retail Banking. The aim of this co-location was to develop synergies between the teaching, research and commercial communities involved in all aspects of database activities, and to use BNCOD as a focus for future synergies and developments within these communities.

Although this is entitled the British National Conference on Databases, BNCOD has always had an international focus, and this year more than most, with the majority of the papers submitted and accepted coming from outwith the UK. We were fortunate in attracting over 70 paper submissions for the conference, of which 18 were accepted, reflecting a stringent and quality-driven peer-review process and maintaining the reputation of BNCOD as a conference reporting high-quality research to the community. However, in our aim to extend the inclusivity of the conference, we also introduced an application paper stream to encourage submissions less focused on research (this stream appears in Vol. 2 of the proceedings), and at the end of this volume you will find the three best papers from TLAD2, as identified by the reviewers for that workshop.

The conference itself reflected the ever-growing diversity of the database community, with papers reflecting traditional database research concerns, through XML to multimedia, but we have strongly resisted the temptation to diversify away from our core concerns. In particular, the importance of database research was reflected in the subjects raised by the two keynote speakers. Both of these subjects are of massive international significance and have achieved the kind of universal importance that permits them to be identified with the definite article – The GRID and The Web.

Our first keynote speaker, Domenico Laforenza, is Technology Director of the Information Science and Technologies Institute (ISTI) of the Italian National Research Council (CNR). An active researcher and promoter of high-performance computing and GRID technologies, he has been deeply involved at both national and European level in the development of GRID infrastructures in Europe and the oversight of future directions for that development. His keynote address “Towards a Next Generation Grid” considered the developments in GRID research and technology that have brought us to the present situation, and elaborated on Next Generation Grid(s) in 2005–2010, and the research opportunities that will result from them, with a particular emphasis on the convergence of multi-disciplinary research.

Our second keynote speaker, Michael Wilson, is Manager of the W3C Office in the UK and Ireland, and a Member of the EU IST programme Advisory Group (ISTAG). His role as Manager in W3C is to help achieve the goal of leading the Web to full potential through ensuring the interoperability of different proprietary systems. As an active researcher in knowledge engineering, HCI, multimedia and VR, he has the background and pedigree to undertake this daunting and complex task. His keynote address “The Future of the World Wide Web” brought together a number of key concepts for this conference, the use of Web services and the Semantic Web, and the relationship between Web and GRID technologies, and he also introduced the issue of trust as a key concept in the future development of all these services.

As identified earlier, we were fortunate in attracting high-quality papers across a range of related topics. The first set of contributed papers were concerned with the processing of queries applied to data streams. Qingchun Jiang and Sharma Chakravarthy introduced novel scheduling strategies to minimize tuple latency and total memory requirements, and presented experimental results on their efficiency and effectiveness. Dan Olteanu, Tim Furche and Francois Bry were concerned with an evaluator (SPEX) for querying XML streams for which the complexity results are polynomial in the sizes of the data and queries, compared to most other methods, which are exponential in the size of the queries.

In the area of integrating data from a set of heterogeneous databases, one problem that needs to be dealt with is that of multiple query languages. Damir Becarevic and Mark Roantree described the EQL language for querying OO and object-relational database schemas in a database- and platform-independent manner. A system that is being developed to perform data integration based on a schema transformation approach is called AutoMed; Lucas Zamboulis described how the integration of XML data sources is accomplished through graph restructuring of their schemas in AutoMed. Kajal Claypool and Elke Rundensteiner focused on the problem of change over time in a database and its effect on the mapping of one data model to another.

In data analytics/manipulations one type of query that would be very useful for data mining is the path query, which determines the relationships between entities in a database. Rachel Hamill and Nigel Martin described one way of extending conventional database technology to support this type of query. Weifeng Chen and Kevin Lu discussed how agent technology has been integrated with conventional data mining techniques to realize a system for financial data mining. In order to limit the damage caused by malicious attacks on a database, Indrakshi Ray et al. considered two techniques for rapid damage assessment and proposed a new one that improves on them.

XML was the focus of much of the work here. With it there was interest in new query languages and query processing techniques. G. Subramanyam and P. Sreenivasa Kumar presented a technique for multiple structural joins and showed that it performs better than existing join algorithms. Jung Kee Park and Hyunchul Kang were concerned with the use of caches to provide efficient support for XML queries, and focused on the problem of XML updates and the effect on cache answerability. Alison Cawsey et al. discussed the use of transformation constraints placed on XML documents by the information provider to constrain transformations applied to it, e.g., in personalization. On a different tack Richard Wheeldon, Mark

Levene and Kevin Keenoy were concerned with keyword search in relational databases and presented a novel algorithm for join discovery.

At the interface with the user, the automatic generation of data entry interfaces to databases usually follows a simplistic approach with limited constraints on data captured. However, Alan Cannon et al. described a semi-automatic tool for generating interfaces, which uses a domain ontology to reflect the semantics of the data, and improve the quality of the captured data. Linas Bukauskas and Michel Bohlen addressed the problem of scalability of visualization systems for large databases when data is extracted from a database and stored in a scene tree. They introduced two new data structures that improve scalability by eliminating the data bottleneck. Images may be stored in a multimedia database as sequences of editing operations. Leonard Brown and Le Gruenwald presented algorithms for performing color-based similarity searches on sets of images stored in this way.

The final session was on spatial databases and there were three papers on this theme. The performance of the M-tree, a data structure to support access to spatial data, depends on the degree of overlap between spatial regions. Alan Sexton and Richard Swinbank presented a new bulk loading algorithm to improve the performance of the M-tree and introduce a variant of it called the SM-tree. Another spatial data problem, the Obstructed Nearest Neighbour, is concerned with finding the nearest neighbours to a point in the presence of obstacles. Chengyi Xia, David Hsu and Anthony Tung presented an efficient algorithm for solving this problem. Despite much research on spatio-temporal data types, data models and query languages, little work has been done on complete spatio-temporal database systems. Tony Griffiths et al. addressed the latter problem with work on support for database programming in the spatio-temporal OODBMS Tripod.

Acknowledgements

We are deeply indebted to a number of different people and organizations without whom BNCOD21 would not have been possible. In particular we would like to thank:

- The members of the programme committee who gave willingly of their time and expertise and were most helpful, especially in having to deal with more papers than anticipated, and in managing to keep to the tight deadlines imposed on us.
- The invited speakers and authors who presented papers at BNCOD as well as those who submitted papers but were not successful. Without you there would have been no conference.
- The sponsors and exhibitors at the workshops and conferences during the whole week; your contributions, be they financial or in time and service, helped to enhance the quality of the experience for the delegates.
- The organizing committee who worked hard behind the scenes to ensure the success of the conference.

Conference Committees

Organizing Committee

L.M. MacKinnon	(Chair)	Heriot-Watt University
A.G. Burger	(BioInformatics Workshop)	Heriot-Watt University
E. Dempster	(Logistics)	Heriot-Watt University
R.G. Dewar	(Publicity, Website)	Heriot-Watt University
P.M.D. Gray	(Competition)	Aberdeen University
D. Lanc	(Data Financials Workshop)	Heriot-Watt University
P.W.Trinder	(Posters, Tutorials)	Heriot-Watt University

Programme Committee

Howard Williams	(Chair)	Heriot-Watt University
David Bell		Queen's University, Belfast
Albert Burger		Heriot-Watt University
Richard Connor		Strathclyde University
Richard Cooper		University of Glasgow
Barry Eaglestone		University of Sheffield
Wenfei Fan		University of Edinburgh
Alvaro Fernandes		University of Manchester
Alex Gray		University of Wales, Cardiff
Peter Gray		University of Aberdeen
Mike Jackson		University of Wolverhampton
Anne James		Coventry University
Keith Jeffery		CLRC, Rutherford Appleton Laboratory
Jessie Kennedy		Napier University
Brian Lings		University of Exeter
Lachlan MacKinnon		Heriot-Watt University
Nigel Martin		Birkbeck College, University of London
Ken Moody		University of Cambridge
Werner Nutt		Heriot-Watt University
Norman Paton		University of Manchester
Alex Poulouvassilis		Birkbeck College, University of London
Stratis Viglas		University of Edinburgh
Peter Wood		Birkbeck College, University of London

Steering Committee

Brian Lings (Chair)

Barry Eaglestone

Alex Gray

Anne James

Keith Jeffery

Roger Johnson

University of Exeter

University of Sheffield

University of Wales, Cardiff

Coventry University

CLRC, Rutherford Appleton Laboratory

Birkbeck College, University of London

Table of Contents

Invited Papers

Towards a Next Generation Grid	1
<i>Domenico Laforenza</i>	
The Future of the World Wide Web?	4
<i>Michael Wilson, Brian Matthews</i>	

Data Streams

Scheduling Strategies for Processing Continuous Queries over Streams	16
<i>Qingchun Jiang, Sharma Chakravarthy</i>	
Evaluating Complex Queries Against XML Streams with Polynomial Combined Complexity	31
<i>Dan Olteanu, Tim Furche, François Bry</i>	

Integration and Heterogeneity

The EGTV Query Language	45
<i>Damir Bećarević, Mark Roantree</i>	
XML Data Integration by Graph Restructuring	57
<i>Lucas Zamboulis</i>	
AUP: Adaptive Change Propagation Across Data Model Boundaries	72
<i>Kajal Claypool, Elke A. Rundensteiner</i>	

Data Analytics/Manipulations

Database Support for Path Query Functions	84
<i>Rachel Hamill, Nigel Martin</i>	
Agent-Based Data Mining Approach for the Prediction of UK and Irish Companies Performance	100
<i>Weifeng Chen, Kevin Lü</i>	
Reducing Damage Assessment Latency in Survivable Databases	106
<i>Indrakshi Ray, Ross M. McConnell, Monte Lunacek, Vijay Kumar</i>	

XML

Efficient Processing of Multiple Structural Join Queries	112
<i>G.V. Subramanyam, P. Sreenivasa Kumar</i>	

Handling Updates for Cache-Answerability of XML Queries
on the Web 124
Jung Kee Park, Hyunchul Kang

Constraining XML Transformations
for Personalised Information Presentation 136
*Alison Cawsey, Euan Dempster, Daniel Pacey, Howard Williams,
David Marwick, Lachlan MacKinnon*

DbSurfer: A Search and Navigation Tool for Relational Databases 144
Richard Wheeldon, Mark Levene, Kevin Keenoy

Interface/Visualization

Ontology-Driven Automated Generation of Data Entry Interfaces
to Databases 150
Alan Cannon, Jessie B. Kennedy, Trevor Paterson, Mark F. Watson

Incremental Observer Relative Data Extraction 165
Linas Bukauskas, Michael H. Böhlen

Performing Colour-Based Similarity Searches in Multimedia
Database Management Systems Augmented with Derived Images 178
Leonard Brown, Le Gruenwald

Spatial Data

Bulk Loading the M-Tree to Enhance Query Performance 190
Alan P. Sexton, Richard Swinbank

A Fast Filter for Obstructed Nearest Neighbor Queries 203
Chenyi Xia, David Hsu, Anthony K.H. Tung

Language Bindings for Spatio-Temporal Database Programming in Tripod 216
*Tony Griffiths, Norman W. Paton, Alvaro A.A. Fernandes,
Seung-Hyun Jeong, Nassima Djafri*

TLAD Workshop

Gauging Students' Understanding Through Interactive Lectures 234
Helen C. Purchase, Christopher Mitchell, Iadh Ounis

Teaching SQL — Which Pedagogical Horse for This Course? 244
Karen Renaud, Judy van Biljon

Computer-Based Assessment Strategies in the Teaching of Databases
at Honours Degree Level 1 257
Gill Harrison

Author Index 265

Towards a Next Generation Grid

Domenico Laforenza

High Performance Computing Laboratory, ISTI, CNR, Via G. Moruzzi 1,
56126 Pisa, Italy
domenico.laforenza@isti.cnr.it

Extended Abstract

This paper aims to present the outcome of a group of independent experts convened by the European Commission with the objective to identify potential European Research priorities for Next Generation Grid(s) in 2005 – 2010 [1].

The first part will be focused on the Grid Evolution. In fact, in order to discuss “what is a Next Generation Grid”, it is important to determine “with respect to what”. Distinct phases in the evolution of Grids are observable. At the beginning of the 90’s, in order to tackle huge scientific problems, in several important research centres tests were conducted on the cooperative use of geographically distributed resources, conceived as a single powerful computer. In 1992, Charlie Catlett and Larry Smarr coined the term “Metacomputing” to describe this innovative computational approach [2].

The term Grid Computing [3] was introduced by Foster and Kesselman a few years later, and in the meanwhile several other words were used to describe this new computational approach, such as Heterogeneous Computing, Networked Virtual Supercomputing, Heterogeneous Supercomputing, Seamless Computing, etc., Metacomputing could be considered as the 1st generation of Grid Computing, some kind of “proto-Grid”.

The Second Grid Computing generation starts around 2001, when Foster et al. proposed Grid Computing as “*an important new field, distinguished from conventional distributed computing by its focus on large-scale resource sharing, innovative applications, and, in some cases, high-performance orientation*” [4].

A Grid provides an abstraction for resource sharing and collaboration across multiple administrative domains. The term resource covers a wide range of concepts including physical resources (computation, communication, storage), informational resources (databases, archives, instruments), individuals (people and the expertise they represent), capabilities (software packages, brokering and scheduling services) and frameworks for access and control of these resources.

With the advent of multiple different Grid technologies the creativity of the research community was further stimulated, and several Grid projects were proposed worldwide. But soon a new question about how to guarantee interoperability among Grids was raised. In fact, the Grid Community, mainly created around the Global Grid Forum (GGF) [5], perceived the real risk that the far-reaching vision offered by Grid Computing could be obscured by the lack of interoperability standards among the cur-

rent Grid technologies. Interoperability is paramount for problem solving; in fact, there is no single Grid environment because data sources, software sources, compute and storage resources, network resources and detection equipment resources are heterogeneous and distributed by nature.

The marriage of the Web Services technologies [6] with the Second Generation Grid technology led to the valuable GGF Open Grid Services Architecture (OGSA) [7], and to the creation of the Grid Service concept and specification (Open Grid Service Infrastructure - OGSI). OGSA can be considered the milestone architecture to build Third Generation Grids. Although not new, an important ongoing process in the Grid evolution is the concept of “virtualization” of Grid resources. Virtualization is a way to expose only high-level functionalities in a standard and interoperable way hiding the implementation details.

The second part of this paper is focused on the “Next Generation Grids”. Starting from the Expert Group’s opinion that current Grid implementations lack many essential capabilities, which would enable the vision of complete resource virtualization, NGG is seen as something that goes over the above mentioned Third Generation Grids. The Next Generation Grid Properties (“The NGG Wish List”) will be presented. The current Grid implementations do not individually possess all of these properties. However, future Grids not possessing them are unlikely to be of significant use and, therefore, inadequate from both research and commercial perspectives.

As reported in [1], NGG will be *“a virtual, pervasive organisation with specific computational semantics. It performs a computation, solves a problem, or provides service to one single client or to millions of clients”*. NGG will consist of millions of interconnected nodes and will pervade into everyday life.

Nowadays, the Grid programmers have to mastermind the usage of different kinds of geographically dispersed resources; all the intricacies related to resource allocation and scheduling, data movement, synchronisation, error handling, load balancing, etc. must be transparent to users and developers. Future generations Grids should be programmed through generic and problem-specific abstractions, supported by an appropriate programming environment. This requires both the needs to study and adapt existing programming models to the Grid context, as well as the definition of new programming models, combining parallel and distributed programming practices in a coherent way.

In order to realise the NGG vision much research is needed. During the last few years, several new terms such as Global Computing, Ubiquitous Computing, Utility Computing, Pervasive Computing, On-demand Computing, Autonomic Computing, Ambient Intelligence [8], etc., have been coined. In some cases, these terms describe very similar computational approaches. The term Ambient Intelligence (AmI), for instance, stems from the convergence of three key technologies: Ubiquitous Computing, Ubiquitous Communication, and Intelligent User Friendly Interfaces. It is not easy in much of the general literature on technology futures to see clearly how AmI is distinguished from older concepts such as Pervasive Computing or Ubiquitous Computing, the latter being first described more than a decade ago.

Consequently, some people are raising the following questions: Are these computational approaches facets of the same medal? What is their relationship with Grid? Moreover, for each of these approaches, we are assisting with the creation of quite disjoint “schools of thought”, having very few cultural interactions with each other.

Often these communities use different languages and terminology just to emphasize the importance of their own technical contributions. A great deal of disagreement exists as to the individual definitions of each of the terms. Sometimes the same term is used to indicate different things, or different terms are used to indicate the same thing. This situation is not necessarily problematic “per se”, although it may be somewhat confusing for people who are looking for more precise definitions of each term. In order to reduce the risk to re-invent the wheel, terms, concepts and definitions need to be standardized.

The realisation of the NGG vision requires much research (and funds), and the real problem that could derive from a persistent situation of apparently confusing and unclarified overlapping of those research areas is the fact that there could be some concrete risks to waste intellectual and physical (e.g. public and private funds) resources in mere “religious” wars. This will not be beneficial for anybody.

In conclusion, paraphrasing Charles Darwin, the idea that each species has been independently created is erroneous, and a stronger convergence of R&D actions among those exciting research fields seems to be absolutely required.

References

1. EU Expert Group Report, Next Generation Grid(s) 2005 – 2010, Brussels, June 2003, ftp://ftp.cordis.lu/pub/ist/docs/ngg_eg_final.pdf.
2. Smarr, L., Catlett, C.: Metacomputing, *Communications of the ACM* **35**(6) (1992) 45-52
3. Foster, I., Kesselman, C. (Editors), *The Grid: Blueprint for a Future Computing Infrastructure*, Morgan Kaufmann Publishers, USA (1999)
4. Foster, I., Kesselman, C., Tuecke, S.: The Anatomy of the Grid: Enabling Scalable Virtual Organizations, *Int. Journal of High Performance Computing Applications* **15**(3) (2001) 200–222
5. Global Grid Forum, www.ggf.org
6. W3C Consortium. www.w3c.org
7. Foster, I., Kesselman, C., Nick, J., Tuecke, S.: *The Physiology of the Grid: An Open Grid Services Architecture for Distributed Systems Integration* (2002)
8. Aarts, E., Bourlard, H., Burgelman, J-C., Filip, F., Ferrate, G., Hermenegildo, M., Hvannberg, E., McAra-McWilliam, I., Langer, J., Lagasse, P., Mehring, P., Nicolai, A., Spirakis, P., Svendsen, B., Uusitalo, M., Van Rijsbergen, K., Ayre, J.: *Ambient Intelligence: from vision to reality*. IST Advisory Group (ISTAG) in FP6: WG116/9/2003 (2003) Final Report.

The Future of the World Wide Web?

Michael Wilson and Brian Matthews

W3C Office in the UK and Ireland
CCLRC Rutherford Appleton Laboratory,
Chilton, Didcot, OX11 0QX, UK
{M.D.Wilson, B.M.Matthews}@rl.ac.uk
<http://www.w3c.rl.ac.uk>

Abstract. The Web started as a simple and very usable distributed system that was rapidly adopted. The Web protocols then passed through a period of rationalization and development to separate content from presentation in order to promote the re-usability of content on different devices. Today the developments in Web technologies are addressing new opportunities in Web Services and the Semantic Web, as well as the growing cultural diversity of the Web. These developments unite in the issue of *trust*, of content and services available on the Web, but also in access by others to the content and services that users may own. While the Web has been rationalizing, the Grid has developed to provide academic science with easier access to services and content. The Grid is now moving to exploit the robust interoperable commodity Web Services instead of maintaining its own middle level infrastructure. As Web Services, the Grid and the Semantic Web develop they will become increasingly interdependent on each other, and indistinguishable from the mainstream Web.

1 The Past

In 1991 Tim Berners-Lee, Robert Cailliau and Nicola Pellow from CERN released a portable line mode browser which could access documents held on distributed servers written in the HyperText Mark-up Language (HTML), through the HyperText Transport Protocol (HTTP), FTP or other protocols from a single address space within which each had a unique Universal Resource Location (URL).

The first major change to HTML came in 1993 when Marc Andreessen and Eric Bina from NCSA wrote the Mosaic browser and allowed in-line colour images through the introduction of the tag.

A major addition to the overall Web architecture was also made in 1993 when Matthew Gray at MIT developed his World Wide Web Wanderer which was the first robot on the Web designed to count and index the Web servers. Initially, it only counted the available Web servers, but shortly after its introduction, it started to capture URLs as it went along. The database of captured URLs became the first Web database - the Wandex.

By 1993 these major components of the Web architecture were available from research organisations. The subsequent growth in Web pages and servers, and the

development of commercial tools, which will not be recounted here. Such growth and commercial involvement show that the Web was a practical success that was becoming prey to competitive commercial interests. In 1994, as a result of concerns that fragmentation of Web standards would destroy the interoperability that the Web had achieved, Tim Berners-Lee and the Laboratory for Computer Science of MIT started the World Wide Web Consortium (W3C) to direct future developments for the Web.

The technologies which constitute the Web, although practical and successful, were technically crude. HTML combined the description of the content of a page with a description of the presentation of the page, which limited the re-usability of the content, and included no type system to support static checking. There was no mechanism to present or compose time-based media such as sound or video. The HTTP transport protocol was not optimized to the resource transfer usage of the Web. There was even confusion about what a URL could point to – files, but what of other resources such as devices (e.g. printers) and even people. However, these limitations allowed an accessible conceptual model and easy interaction without which it is unlikely that the Web would have been adopted.

W3C entered into a programme of work to reform these protocols, to overcome these problems, and incorporate technologies that would facilitate extensibility. To do this the simple architecture of URL, HTML and HTTP alone had to be sacrificed for a more complex one, where the new protocols and languages would no longer be easy to write in generic text editors, but would require specialized editors.

W3C addresses the core Web technologies that build on the transport layer standardized by the IETF and which are, in turn, built on by application specific standards that require a less rigorous process. Once a W3C working group is established to create a standard, or recommendation as W3C calls them, it usually takes two to three years before it is completed and officially published. Some groups have been terminated before completion when the motivation for standardization has dissipated or stakeholder consensus cannot be reached. The protocols published so far as recommendations by W3C are shown in Table 1.

2 The Present

The main four concerns of W3C today are listed below, while the architecture as it has evolved to meet them is shown in Figure 1.

Ensure access to the Web by many devices – The Web is becoming accessible from a wide range of devices including cellular phones, TV, digital cameras, and in-car computers. Interaction with resources on the Web can be achieved through a key pad, mouse, voice, stylus or other input devices. W3C has activities addressing device independence and multimodal interaction to contribute to W3C's goal of universal access.

Account for cultural diversity – To ensure access to the Web by people speaking different languages, with different writing conventions, and having different cultural backgrounds. In 1999, approximately half of all public Web sites were associated

with entities located in the United States, whereas by March 2003 only 35.2% of Web pages were in English, with 22.2% in Chinese or Japanese.

Web Services – the Web started as a tool for users to view documents. It is now moving to be a tool for computers to communicate with each other, on services provided by each other in a peer-to-peer manner rather than only in a browser to server one.

Table 1. Recommendations issued by W3C before March 2004

1996	1997	1998	1999	2000	2001	2002	2003	2004
PICS Rating	PICSRules	PICS DSig	Name spaces	DOM 2 Core	XHTML M12n	P3P 1.0	SVG 1.1	DOM 3 Validation
PICS Labels	HTML 3.2	CSS 2 SMIL 1.0 DOM 1	CSS 1 WCAG 1.0 Style Sheets PI	ATAG 1.0 DOM 2 Events DOM 2 Style	Canonical XML Schema Primer Schema Struct.	XML Signature XML Canonicalization XHTML 1.0	DOM 2 HTML SVG Mobile XPTR Element	CC/PP Infoset (2nd) Namespaces 1.1
			MathML 1.01 XPath 1.0 XSLT 1.0 HTML 4.01	DOM2T raversal DOM 2Views XHTML Basic	Schema Types Ruby XHTML 1.1 XLink 1.0 XML Base SMIL 2.0 SMIL Anim. XSL 1.0 WebCGM	XPath Filter Decrypt Transform XML Encryption UAAG1.0	XPTR Framework XPTR Xmlns SOAP Adjuncts SOAP Framework SOAP Primer SOAP Tests XForms 1.0 XML Events MathML 2.0 PNG (2nd)	XML 1.0 (3rd) XML 1.1 OWL Guide OWL Overview OWL Reference OWL Semantics OWL Tests OWL Use Cases RDF Concepts RDF Primer RDF Schema RDF Semantics RDF Test Cases RDF/XML Speech Recognition VoiceXML 2.0

Semantic Web – as the data and service servers become more common on the Web, the data and services that they provide need to be described in a machine understandable way in order to be discovered, evaluated for fitness to a user's purpose

and then called. The Web also needs to provide an environment in which contracts for the use of data and services can be established, and trust relationships defined to limit the growing incidence of cybercrime. To support machine dialogues in these areas, richer representations of ontologies, rules, and inference are required which are collectively termed the Semantic Web.

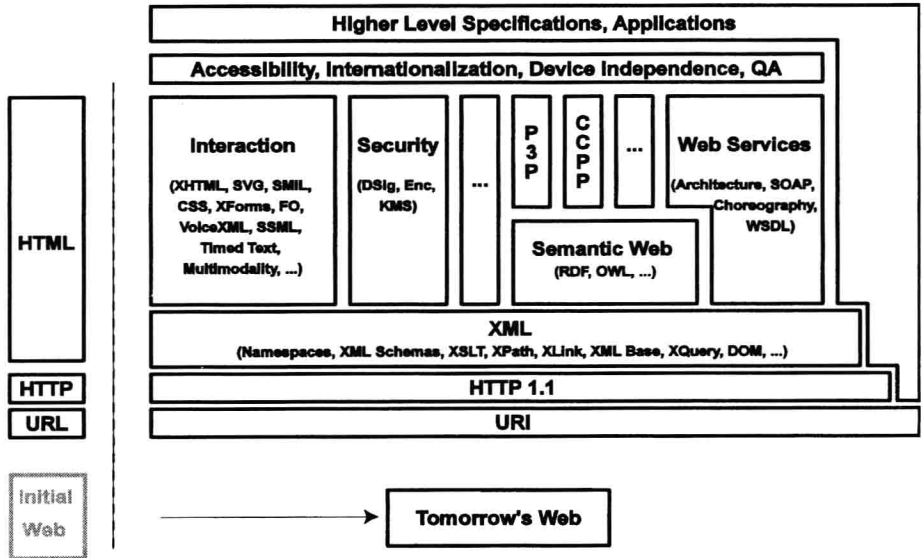


Fig. 1. The evolution of the architecture of the protocols in the World Wide Web

3 Web Services

Web Services have a long ancestry in distributed computing going back to remote procedure calls. The XML Protocol Activity which became the Web Services protocol or SOAP layer in the Web Services architecture was initiated in September 2000 in W3C following the observations that “distributed object oriented systems such as CORBA, DCOM and RMI exist with distinct functionality and distinct from the Web address space causes a certain tension, counter to the concept of a single space”[1]. As shown in Table 1, it was 2003 before any parts of SOAP reached the final recommendation form of publication.

In 2002 IBM and Microsoft agreed the main structure of the Web Services Architecture shown in Figure 2 (after [2]) which incorporated the Web Services Security component that is enlarged in Figure 3 (after [3]) since it has been so subdivided. These figures show the overall approach of a transport layer (HTTP) carrying messages (SOAP) between machines to invoke services. These services must describe their functionality (WSDL) and such descriptions can be registered in a directory (UDDI) to ease resource discovery. Privacy issues are still awaiting