GWDG-Bericht Nr. 62

Kurt Kremer, Volker Macho (Hrsg.)

Forschung und wissenschaftliches Rechnen

Beiträge zum Heinz-Billing-Preis 2002

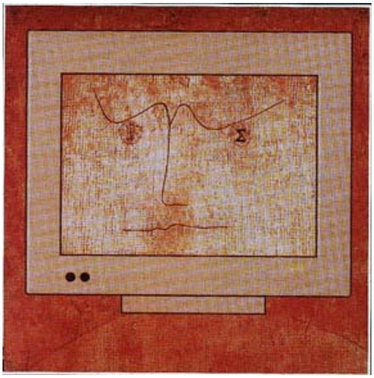

Gesellschaft für wissenschaftliche Datenverarbeitung mbH Göttingen

Forschung und
wissenschaftliches Rechnen

Titelbild:
Logo nach Paul Klee „Der Gelehrte", Modifizierung durch I. Tarim,
Max-Planck-Institut für Psycholinguistik, Nijmegen.

GWDG-Bericht Nr. 62

Kurt Kremer, Volker Macho (Hrsg.)

Forschung und wissenschaftliches Rechnen

Beiträge zum Heinz-Billing-Preis 2002

Gesellschaft für wissenschaftliche Datenverarbeitung
Göttingen

2003

Die Gesellschaft für wissenschaftliche Datenverarbeitung mbH Göttingen ist eine gemeinsame Einrichtung der Max-Planck-Gesellschaft zur Förderung der Wissenschaften e. V. und des Landes Niedersachsen. Die Max-Planck-Gesellschaft hat diesen Bericht finanziell gefördert.

Redaktionelle Betreuung:
Volker Macho (Max-Planck-Institut fuer Polymerforschung, Mainz)

Satz:
Gesellschaft für wissenschaftliche Datenverarbeitung mbH Göttingen,
Am Faßberg, D-37077 Göttingen

Druck: klartext GmbH, D-37073 Göttingen, www.kopie.de

ISSN: 0176-2516

Inhalt

Vorwort

Der vorliegende zehnte Band der Reihe „Forschung und wissenschaftliches Rechnen“ enthält fünf der sechs für den Heinz-Billing-Preis im Jahre 2002 eingereichten Beiträge. Ein Beitrag lag leider bis zur Drucklegung nicht vor. Die diesjährigen Preisträger waren Daan Broeder, Hennie Brugman und Reiner Dirksmeyer vom Max-Planck-Institut für Psycholinguistik in Nijmegen. Sie wurden ausgezeichnet für ihre Arbeit „NILE“ (Nijmegen Language Resource Environment), ein international viel beachtetes Werkzeug zur computer-basierten Verarbeitung multimedialer Informationen.

Erstmalig wurde hiermit der Heinz-Billing-Preis für ein wissenschaftliches Projekt vergeben, welches nicht im Bereich der klassischen Naturwissenschaften angesiedelt ist. Zwei weitere eingereichte Beiträge, nämlich „Beitrag zum E-Learning“ von H. Ibelgaufts sowie „The Virtual Laboratory Infrastructure for Controlled Experiments in Economics“ von B. Greiner et. al. fallen ebenfalls in diese Kategorie.

Auch die übrigen Beiträge aus den Bereichen Physik und Medizin, „Gene“, ein Programm zur Untersuchung turbulenter Transportvorgänge in Fusionsplasmen von F. Jenko und „Introducing Lambda Tensor 1.0“ von Th. Fischbacher sowie die Arbeit „Trax“, ein Programmpaket zur Visualisierung von Nervenfasern in MRI-Daten von B. Pütz, K. Ates und D. Auer, zeigen das hohe Niveau wissenschaftlicher Datenverarbeitung.

Während in den letzten Jahren die Ausschreibung des Preises im wesentlichen auf die Institute der Max-Planck-Gesellschaft beschränkt war, hat die letzte Mitgliederversammlung beschlossen, die Ausschreibung zu öffnen und auf die Universitäten des Landes auszudehnen. Wir versprechen uns damit eine noch größere Breite der Beiträge, die dann auch die große Vielfalt der DV-Anwendungen in der Wissenschaft widerspiegeln.

An dieser Stelle möchten wir insbesondere Prof. H. Billing, dem Stifter des Preises danken, dass er dieser Neuerung zugestimmt und sie befürwortet hat.

Ein ganz besonderer Dank sei an dieser Stelle Herrn Günter Koch, GWDG, für die Umsetzung der eingesandten Manuskripte in eine für das Offsetdruckverfahren kompatiblen Druckvorlage ausgesprochen.

Die Vergabe des Preises in dieser Form wäre ohne einen Hauptsponsor nicht möglich. Wir danken der Fa. IBM Deutschland, die dies auch in diesem Jahr wieder übernommen hat.

Beachten Sie bitte, dass sich die Adresse unseres Web-Auftritts geändert hat. Informationen zum Heinz-Billing-Preis und zu den Arbeiten der letzten Jahre sind nunmehr unter

www. billingpreis.mpg.de

im Internet zu finden.

Kurt Kremer, Volker Macho

Der Heinz-Billing-Preis 2002

Ausschreibung des Heinz-Billing-Preises 2002 zur Förderung des wissenschaftlichen Rechnens

Im Jahre 1993 wurde zum ersten Mal der Heinz-Billing-Preis zur Förderung des wissenschaftlichen Rechnens vergeben. Mit dem Preis sollen die Leistungen derjenigen anerkannt werden, die in zeitintensiver und kreativer Arbeit die notwendige Hard- und Software entwickeln, die heute für neue Vorstöße in der Wissenschaft unverzichtbar sind.

Der Preis ist benannt nach Professor Heinz Billing, emeritiertes wissenschaftliches Mitglied des Max-Planck-Institutes für Astrophysik und langjähriger Vorsitzender des Beratenden Ausschusses für Rechenanlagen in der Max-Planck-Gesellschaft. Professor Billing stand mit der Erfindung des Trommelspeichers und dem Bau der Rechner G1, G2, G3 als Pionier der elektronischen Datenverarbeitung am Beginn des wissenschaftlichen Rechnens.

Der Heinz-Billing-Preis zur Förderung des wissenschaftlichen Rechnens steht unter dem Leitmotiv

„EDV als Werkzeug der Wissenschaft“.

Es können Arbeiten eingereicht werden, die beispielhaft dafür sind, wie die EDV als methodisches Werkzeug Forschungsgebiete unterstützt oder einen neuen Forschungsansatz ermöglicht hat.

Der folgende Stichwortkatalog mag den möglichen Themenbereich beispielhaft erläutern:

– Implementation von Algorithmen und Softwarebibliotheken
– Modellbildung und Computersimulation
– Gestaltung des Benutzerinterfaces
– EDV gestützte Meßverfahren
– Datenanalyse und Auswertungsverfahren
– Visualisierung von Daten und Prozessen

Die eingereichten Arbeiten werden referiert und in der Buchreihe "Forschung und wissenschaftliches Rechnen" veröffentlicht.

Die Jury wählt einen Beitrag für den mit € 3000,- dotierten Heinz-Billing-Preis 2002 zur Förderung des wissenschaftlichen Rechnens aus. Die Beiträge, in deutscher oder englischer Sprache abgefasst, müssen keine Originalarbeiten sein und sollten möglichst nicht mehr als fünfzehn Seiten umfassen.

Da zur Bewertung eines Beitrages im Sinne des Heinz-Billing-Preises neben der technischen EDV-Lösung insbesondere der Nutzen für das jeweilige Forschungsgebiet herangezogen wird, sollte einer bereits publizierten Arbeit eine kurze Ausführung zu diesem Aspekt beigefügt werden.

Der Heinz-Billing-Preis zur Förderung des wissenschaftlichen Rechnens wird jährlich vergeben. Die Preisverleihung findet anlässlich des 19. DV-Treffens der Max-Planck-Institute am 21. November 2002 in Göttingen statt.

Beiträge für den Heinz-Billing-Preis 2002 sind bis zum 1. September 2002 einzureichen.

Heinz-Billing-Preisträger

1993: Dr. Hans Thomas Janka, Dr. Ewald Müller, Dr. Maximilian Ruffert
Max-Planck-Institut für Astrophysik, Garching

Simulation turbulenter Konvektion in Supernova-Explosionen in massereichen Sternen

1994: Dr. Rainer Goebel
Max-Planck-Institut für Hirnforschung, Frankfurt

- Neurolator - Ein Programm zur Simulation neuronaler Netzwerke

1995: Dr. Ralf Giering
Max-Planck-Institut für Meteorologie, Hamburg

AMC: Ein Werkzeug zum automatischen Differenzieren von Fortran Programmen

1996: Dr. Klaus Heumann
Max-Planck-Institut für Biochemie, Martinsried

Systematische Analyse und Visualisierung kompletter Genome am Beispiel von S. cerevisiae

1997: Dr. Florian Mueller
Max-Planck-Institut für molekulare Genetik, Berlin

ERNA-3D (Editor für RNA-Dreidimensional)

1998: Prof. Dr. Edward Seidel
Max-Planck-Institut für Gravitationsphysik, Albert-Einstein-Institut, Potsdam

Technologies for Collaborative, Large Scale Simulation in Astrophysics and a General Toolkit for solving PDEs in Science and Engineering

1999: Alexander Pukhov
Max-Planck-Institut für Quantenoptik, Garching

High Performance 3D PIC Code VLPL:
Virtual Laser Plasma Lab

2000: Dr. Oliver Kohlbacher
Max-Planck-Institut für Informatik, Saarbrücken
BALL – A Framework for Rapid Application Development in Molecular Modeling

2001: Dr. Jörg Haber
Max-Planck-Institut für Informatik, Saarbrücken
MEDUSA, ein Software-System zur Modellierung und Animation von Gesichtern

Das Kuratorium des Heinz-Billing-Preises

Prof. Dr. Heinz Billing
Emeritiertes Wissenschaftliches Mitglied des Max-Planck-Institut für Astrophysik, Garching

Prof. Dr. Friedel Hossfeld
Forschungszentrum Jülich GmbH, Jülich

Prof. Dr. K. Ulrich Mayer
Max-Planck-Institut für Bildungsforschung, Berlin

Prof. Dr. Stefan Müller
Max-Planck-Institut für Mathematik in den Naturwissenschaften, Leipzig

Prof. Dr. Jürgen Renn
Max-Planck-Institut für Wissenschaftsgeschichte, Berlin

Prof. Dr. H. Wolfgang Spiess
Max-Planck-Institut für Polymerforschung, Mainz

Mitglieder der Jury 2002

Prof. Dr. Wolfgang Baumeister
Max-Planck-Institut für Biochemie, Martinsried

Prof. Dr. Peter Hagoort
F.C. Donders Centre for Cognitive Neuroimaging, University of Nijmegen

Prof. Dr. Friedrich Hertweck
Emeritiertes Wissenschaftliches Mitglied
Max-Planck-Institut für Plasmaphysik, Garching

Prof. Dr. Kurt Kremer
Max-Planck-Institut für Polymerforschung, Mainz

Prof. Dr. Stefan Müller
Max-Planck-Institut für Mathematik in den Naturwissenschaften, Leipzig

Prof. Dr. Walter Thiel
Max-Planck-Institut für Kohlenforschung, Mülheim an der Ruhr

Prof. Dr. Simon D.M. White
Max-Planck-Institut für Astrophysik, Garching

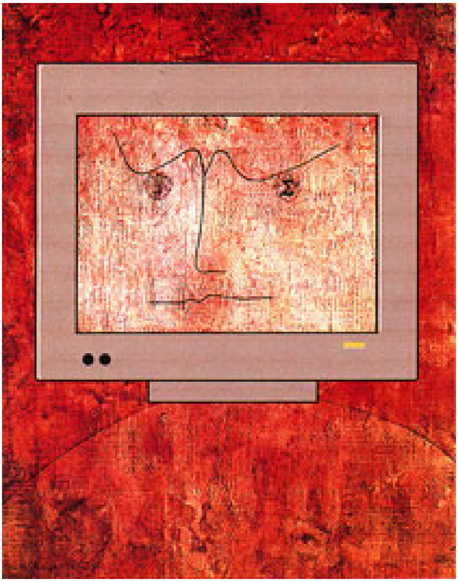

Daan Broeder, Hennie Brugman, Reiner Dirksmeyer
Max-Planck-Institut für Psycholingustik, Nijmegen

erhalten den

Heinz-Billing-Preis 2002
zur Förderung
des wissenschaftlichen Rechnens

als Anerkennung ihrer Arbeit

NILE – Nijmegen Language Resource Environment

Laudatio

Der Heinz-Billing-Preis 2002 wird verliehen für das Programm-paket **NILE** (**Ni**jmegen **L**anguage Resource **E**nvironment). Durch diese Entwicklung wird eine substantielle Veränderung in den Möglichkeiten des wissenschaftlichen Umgangs mit Sprache durch den Einsatz der modernen Informationstechnologie erreicht. Die wissenschaftliche Aufarbeitung eines der größten Korpora mit multimedialen Sprachressourcen, u.a. vieler vom Aussterben be-drohter Sprachen, wird damit wesentlich vereinfacht und diese bes-ser zugänglich gemacht. Das Paket benutzt dabei neueste Standards, die seiner Verbreitung schon jetzt sehr zu Gute kommen. Nile hat dabei auch einen wichtigen Anstoß zur konkreten Aus-füllung des Semantic Web gegeben.

Herr Prof. Kurt Kremer überreicht Herrn Daan Broeder die Urkunde zum Heinz-Billing-Preis 2002

NILE: Nijmegen Language Resource Environment[1]

Daan Broeder, Hennie Brugman, Reiner Dirksmeyer

MPI für Psycholinguistik

1. Einleitung

Am MPI für Psycholinguistik wurde eine umfassende Infrastruktur zum professionellen Umgang mit multimedialen/multimodalen sprachlichen Ressourcen aufgebaut. Diese Infrastruktur deckt den gesamten Lebenszyklus derartiger, in komplexer Weise miteinander verbundener Ressourcen ab, die in der datenorientierten Linguistik und vieler angrenzender Forschungsgebiete[2] von zentraler Bedeutung sind. Daher ist die Verfügbarkeit einer professionellen Umgebung zum Erzeugen, zum Anreichern, zur Bearbeitung, zum Wiederauffinden und zur sicheren Aufbewahrung derartiger Ressourcen für die Forschung essentiell.

Am MPI wurde eine solche Umgebung zunächst vor allem deshalb aufgebaut, um den eigenen Wissenschaftlern ein effizientes Arbeiten mit

[1] Bei NILE handelt es sich um das Arsenal an Methoden und Tools, die im Rahmen der Erstellung und der Pflege eines umfangreichen, der Wissenschaft dienenden multimedialen Sprach-Korpus erarbeitet wurden.

[2] Als Beispiele für angrenzende Forschungsdisziplinen sind hier genannt: Psycholinguistik, Phonetik, Computerlinguistik und die digitale Sprachverarbeitung.

sprachlichen Ressourcen zu ermöglichen. Dabei konnten immer wieder neue Konzepte und Werkzeuge vorgestellt und entwickelt werden, die auch im internationalen Maßstab Beachtung fanden. Mehrere Projekte wie ISLE[3], MUMIS[4], DOBES[5], COREX[6] erlaubten es uns, diese Konzepte umzusetzen. Auch die weitere Entwicklung ist durch neue Projekte wie INTERA und ECHO gewährleistet, die allerdings auch neue Anforderungen beinhalten. Um Anschluß an die internationale Entwicklung zu bewahren, wurde eine Mitarbeit in der neuen ISO Fachgruppe ISO-TC37/SC4[7] in den Bereichen „Annotations-Strukturen", „Lexikalische Strukturen" und „Metadaten" begonnen.

Auf der Basis erster Überlegungen und einer ersten Generation von Werkzeugen wurden die Konzepte für die Infrastruktur im Jahre 1998 zum ersten Male auf der LREC 1998 Konferenz vorgestellt. Auf der LREC 2000 Konferenz wurden erste Versionen und Konkretisierungen vorgestellt, die auch die Grundlage für die weitere Arbeit formten. Auf der LREC Konferenz 2002 konnten die Resultate der Arbeit und auch professionelle Werkzeuge demonstriert werden. Auch konnten die bisherigen Arbeiten in einen weiteren Zusammenhang gestellt und damit Perspektiven für die Zukunft aufgezeigt werden.

2. Motivation

Während die Datenverarbeitung in den naturwissenschaftlichen Disziplinen bereits sehr frühzeitig zum Berechnen wissenschaftlicher Probleme eingesetzt wurde, folgte die Linguistik mit ihren angrenzenden Disziplinen[8] erst relativ spät. Ausnahmen waren Disziplinen, die sich z.B. mit Formalen Sprachen und der digitialen Sprachverarbeitung beschäftigten. Wesentliche Gründe sind die Datenintensität der Linguistik und die Schwierigkeiten, Phänomene der natürlichen Sprachen algorithmisch behandeln zu wollen. Eine Veränderung hat hier erst der große Erfolg statistischer Modelle gebracht, der zunächst in der Verarbeitung gesprochener Sprache und später auch in der automatischen Verarbeitung von Texten zu bemerkenswerten Erfolgen führte.

[3] International Standards for Language Engineering, finanziert durch die EU vom Januar 2000 bis zum September 2002; http://www.mpi.nl/ISLE

[4] Multimedia Indexing and Searching, finanziert durch die EU vom Juli 2000 bis zum Dezember 2003; http://parlevink.cs.utwente.nl/projects/mumis/

[5] Dokumentation bedrohter Sprachen, finanziert durch die VolkswagenStiftung vom September 2000; http://www.mpi.nl/DOBES

[6] Corpus Exploitation tool, finanziert durch NWO (NL) im Rahmen des Spoken Dutch Corpus Projektes; http:// lands.let.kun.nl/cgn/home.htm

[7] ISO TC37/SC4 -erminology and Management of Language Resources, gegründet im May 2002; http://kibs.kaist.ac.kr/iso_tc37_sc4/main.htm

[8] Im folgenden wird der Term „Linguistik" in seiner breiteren Bedeutung verwendet.

Diese Art von Arbeiten waren jedoch zunächst auf wenige spezialisierte Institute beschränkt. Insbesondere jedoch das Aufkommen der Personal Computer erlaubte es einer größeren Zahl von Wissenschaftlern auch in der Linguistik, den Computer für ihre datenorientierte Analysearbeit einzusetzen. Insbesondere wegen der Speicherplatzbeschränkungen etablierte sich für viele Jahre die folgende Vorgehensweise:

1. Entsprechend der wissenschaftlichen Fragestellung wurden zumeist Audio-Aufnahmen gemacht.
2. Anhand eines Transkriptions-Systems wurden die Aufnahmen in computer-basierte Texte überführt, die dann im folgenden als Basismaterial verwendet wurden, d.h. die ursprünglichen Aufnahmen wurden nur in den seltensten Fällen wieder zu Rate gezogen.
3. Der Umfang möglicher Analysen blieb durch die Wahl der Transkribierung beschränkt, d.h. daß zum Beispiel wichtige syntaktische oder semantische Disambiguierungen nicht vorgenommen werden konnten, da die hierfür erforderlichen prosodischen Parameter zumeist nicht codiert waren und der Zugriff auf die ursprünglichen Aufnahmen faktisch nicht realisierbar war.

Hinzu kam, daß immer mehr Forscher die Relevanz von Videoaufnahmen erkannten, um die verbale Kommunikation in ihren multimodalen Kontext zu setzen und um zum Beispiel die Sprache anderer Kulturen auf ihrem kulturellen und soziologischen Hintergrund zu dokumentieren. Ebenfalls im Bereich des „Language Engineering“ erkannte man die Relevanz multimodaler Kommunikation, um die bestehenden Beschränkungen von Mensch-Maschine-Schnittstellen zu überwinden.

Am MPI für Psycholinguistik, als einem der führenden Institute weltweit im Bereich der linguistischen Forschung, wurden diese Trends, wie auch die zunehmenden Probleme mit der traditionellen Arbeitsweise sehr frühzeitig evident. Es wurden daher frühzeitig Möglichkeiten untersucht, um allen Wissenschaftlern und sonstigen Beteiligten einen Online-Zugang zu ihren Materialien zu geben. Der vollständige und konsequente Übergang zu einer computer-basierten Verarbeitung multimedialer Informationen konnte allerdings nur gelingen, wenn alle Aspekte durchdacht und den Wissenschaftlern für alle Situationen attraktive Alternativen angeboten würden. Trotz anfänglicher Zurückhaltung der Wissenschaftler wurde sehr schnell klar, welche enormen Vorteile und welch großes Potential die Digitalisierung mit sich bringt, vorausgesetzt, daß effiziente Verarbeitungsmethoden angeboten und Zusagen zur sicheren Aufbewahrung gemacht werden können.

Gegenwärtig kann weltweit eine extreme Zunahme an (multimedialen/multimodalen) Sprach-Ressourcen festgestellt werden, die die Verfügbarkeit von Infrastrukturen, so wie sie am MPI entwickelt wurden, zu einer Notwendigkeit machen. Im Rahmen der Definition des Begriffes des

Semantic Web[9] wurde deutlich gemacht, daß wir neben dem jetzigen Web, das einem gewaltigen, aber kaum beherrschbaren Hypertext Archiv gleichkommt, unter anderem Metadaten-Beschreibungen der Ressourcen benötigen, um diese in einem professionellen Arbeitsumfeld effizient wiederauffinden zu können, ihre Verwaltung beherrschbar zu machen und sie in Zukunft durch intelligente Agenten bearbeitbar machen zu können. Metadaten, so wie sie im Rahmen von NILE entwickelt wurden, werden einer der Stützpfeiler des Semantic Web sein. Ein Verweis auf andere Projekte zeigt, daß das Grundkonzept an sich nicht neu ist. Die Text Encoding Initiative[10] (TEI) schlug ausführliche Metadaten vor, um Ressourcen zu beschreiben. Verschiedene Projekte, als Beispiel sei hier das CHILDES Projekt[11] genannt, haben Metadaten in Datei-Headern aufgenommen, um sich schnell über deren Inhalt orientieren zu können. Diese Metadaten waren jedoch nicht dazu entworfen, um im Web projekt-unabhängig schnell geeignete Ressourcen finden zu können. Im wesentlichen hat erst die Dublin Core (DC) Initiative[12] diesen Ansatz verfolgt. Allerdings entspringt der DC Metadaten-Satz, der aus 15 vage und breit formulierten Elementen besteht, dem Interesse der Bibliothekare und Archivare nach einem Deskriptoren-Satz, mit dessen Hilfe sich jede Ressource im Web beschreiben läßt. Dieser Satz, der den Ansprüchen der Bibliothekare und eventuell auch des gelegentlichen Web-Benutzers genügen könnte, erfüllt allerdings nicht die Erwartungen der Speziaisten einer Domäne. Für den Bereich sprachlicher Ressourcen wurde daher an einem neuen Ansatz gearbeitet.

Ebenfalls bedarf es effizienter Annotations- und Analyse-Tools für multimediale Resourcen, um die menschliche Interaktion in ihrer Gänze besser verstehen zu können. Im Rahmen des Web gilt es auch hier, Standards wie XML zur Beschreibung der Annotationen einzusetzen und flexible Formate zu definieren, die die erforderliche Komplexität repräsentieren können. Im Rahmen von NILE wurde auch hieran gearbeitet, denn es werden letztlich unter anderem diese Standards sein, die eine Interoperabilität ermöglichen. Bis zum Beginn unserer NILE Arbeiten gab es keine einfach verwendbare und flexibel einsetzbaren multimedialen Annotations-Tools. Das Signstream Tool[13] lief nur auf dem MAC und mit ihm konnten nur beschränkte Arten von Annotationen erzeugt werden. Eine XML-Ausgabe war nicht möglich. Ebenso war das am MPI entwickelte Vorläufer Tool Media-Tagger[14] in seinen Möglichkeiten beschränkt, da das implementierte Datenmodell sehr beschränkt und es ebenfalls nur auf dem MAC ausführbar war.

---

[9] T. Berners-Lee, J. Hendler, O. Lassila. “The Semantic Web”. Scientific American. May 2001

[10] Text Encoding Initiative; http://www.tei-c.org

[11] Child Language Data Exchange System; http://childes.psy.cmu.edu/

[12] Dublin Core; http://dublincore.org

[13] Signstream; http://www.bu.edu/asllrp/SignStream/

[14] Media Tagger; http://www.mpi.nl/world/tg/CAVA/mt/

Erst im ATLAS Modell[15] wurden ähnliche Zielsetzungen wie im NILE Projekt verfolgt. Für beide Projekte lieferte Bird und Liberman's Arbeit am Annotation Graph Modell[16] eine wichtige theoretische Basis.

3. Herausforderungen

Im folgenden sollen die Herausforderungen und die wesentlichen Eckpfeiler der am MPI entwickelten Infrastruktur beschrieben werden.

Der Übergang zu digitalen linguistischen Ressourcen beinhaltet einen Verzicht auf traditionelle Audio/Video-Archive und Audio/Video-Geräte zu deren Speicherung, Präsentation und Analyse. Die digitalen Ressourcen müssen sowohl im Institut als auch im Feld bzw. auf Reisen[17] zur Verfügung stehen, um zu nahezu jedem Zeitpunkt am Material arbeiten zu können.

Da es sich oftmals um wertvolle und nicht wiederholbare Aufnahmen handelt, muß die Speicherung allen Materials entsprechend allgemeiner offener Standards erfolgen und eine größtmögliche Qualität und Orthogonalität der verschiedenen Codierungen angestrebt werden. Ein Konzept für eine Langzeitspeicherung muß entwickelt werden, das sich nicht nur auf ein einzelnes Max-Planck-Institut verlassen darf und das garantiert, daß Links zu den Ressourcen stabil bleiben.

Aufgrund dieser Anforderung (Verteilung der Daten zur sicheren Aufbewahrung) und der im digitalen Zeitalter möglichen und auch geforderten Kollaboration mußten alle Überlegungen und Werkzeuge von einem verteilten Szenario ausgehen, d.h. daß sogar eng zusammengehörende Daten wie zum Beispiel eine Videoaufnahme und die dazugehörige Annotationen der Gestik der aufgenommenen Personen auf verschiedenen Rechnern im Netz gespeichert werden können. Die internationale Kollaboration bekommt durch das Web vollkommen neue Impulse, die den Wunsch, Ressourcen anderer in einfacher Art und Weise wiederzuverwenden unabhängig von derem Speicherort, verstärken wird.

In einem solchen Szenario mit vielfältig relatierten linguistischen Datentypen[18] mußte eine neue Management und Zugriffsebene definiert werden, die den linguistischen Benutzer nicht mit Fragen nach Rechneradressen, den Namen von Plattenlaufwerken und anderen systembedingten Größen belastet, sondern es ihm erlaubt, in einer virtuellen, durch seine

---

[15] ATLAS; http://www.nist.gov/speech/atlas/

[16] Annotation Graph-Modell; http://acl.ldc.upenn.edu/W/W99/W99-0301.pdf

[17] Hierbei kann es sich natürlich nur um einen beschränkten Set von Ressourcen handeln, wenn der Benutzer Notebook nicht ans Interent angschlossen ist.

[18] Mit dem Begriff des Linguistischen Datentyps werden die verschiedenen Typen von Ressourcen bezeichnet als da sind Audio/Video-Aufnahmen, Annotationen zu diesen Aufnahmen, Lexika, Grammatiken, Feld-Notizen, Notizen über das Klangsystem einer Sprache und anderes mehr.

linguistischen Konzepte definierten Welt zu navigieren. Solche typischen linguistischen Konzepte sind der Datentyp, die enthaltenen Sprachen, die verwendeten Modalitäten, das Genre der Aufnahme, das Alter der interviewten Personen und anderes mehr. Eine solche Metadaten-Ebene kann sinnvollerweise nur in einem internationalen Rahmen definiert werden, d.h. die am Institut ausgetesteten Methoden mußten in einen größeren Rahmen eingebettet werden. Ebenfalls galt es, für diese notwendigerweise verteilt existierende Ebene der Metadaten Beschreibungen geeignete Tools zur Verfügung zu stellen.

Als wichtigstem Eckpunkt galt es, den Wissenschaftlern neue Software-Werkzeuge im Umgang mit verteilt vorliegenden multimedialen Ressourcen zu bieten, die sowohl im Netzwerk als auch lokal auf einem Notebook arbeiten. Diese Werkzeuge müssen z.B.

- die effiziente zeitsynchrone Annotation in verschiedenen Sprachen (Fonts, Scriptsysteme) ermöglichen,
- die Exploitation (Suche nach beliebigen Mustern, Visualisierung komplexer Annotationen und deren statistische Auswertung) unterstützen und
- die semi-automatische Interaktion zwischen verschiedenen Datentypen (z.B. Annotationen und Lexikon) bereitstellen.

Angesichts der herrschenden und die wissenschaftliche Arbeit hemmenden Vielfalt der Formate und Strukturen der wesentlichen Datentypen galt es, diese Tools auf abstrakte Daten-Modelle aufzubauen, die ein hohes Maß an Allgemeingültigkeit besitzen. Demzufolge würden es diese Werkzeuge ermöglichen, die verschiedenen existierenden Formate über Import-Module ohne Informationsverlust einzulesen und mittels der für das abstrakte Modell implementierten Operationen zu behandeln.

Zusammen mit den Benutzern mußten neue Workflow Mechanismen erarbeitet und eingeübt werden, da sich die Arbeitsweise der Wissenschaftler mit dem Einzug der digitalisierten Repräsentation aller linguistisch relevanter Daten grundlegend veränderte und langjährig erworbene Fähigkeiten und Fertigkeiten durch neue ersetzt werden mußten. Allgemein wird die Bedeutung dieses Aspektes stark unterschätzt, was leicht zu Akzeptanz-Problemen und damit zu einer Verzögerung der Umsetzung moderner, durch die technologische Entwicklung ermöglichter Konzepte führt.

4. Stand der Arbeiten

In diesem Abschnitt wird dargestellt, wie durch die NILE Umgebung auf die Herausforderungen reagiert wurde, welche Konzepte und Werkzeuge entwickelt wurden und wie weit der gegenwärtige Stand der Entwicklungen gediehen ist.

Die zwei folgenden Bilder zeigen die kanonische Form der Entstehung (Abb. 1) bzw. der Annotation und Analyse (Abb. 2) der am MPI und anderweitig existierenden multimedialen Ressourcen. Die wesentlichen Komponenten, der Browsable Corpus Tool Set (BC) und das Annotation und Analyse Tool ELAN, erlauben vielfältigere Interaktionen und mehrere Datenflows als in dieser Darstellung aufgenommen werden konnte. Abb. 1 zeigt die wesentlichen Pfade zur Erzeugung des Korpus am MPI. Wichtig ist dabei, daß das Workflow darauf ausgerichtet ist, daß jeweils sofort Metadaten Beschreibungen erzeugt werden, wenn die digitale Ressource generiert wird. Nur so kann sichergestellt werden, daß das Archiv von Beginn an wartbar ist und die Benutzer sofort auf die integrierte Resource zugreifen können.

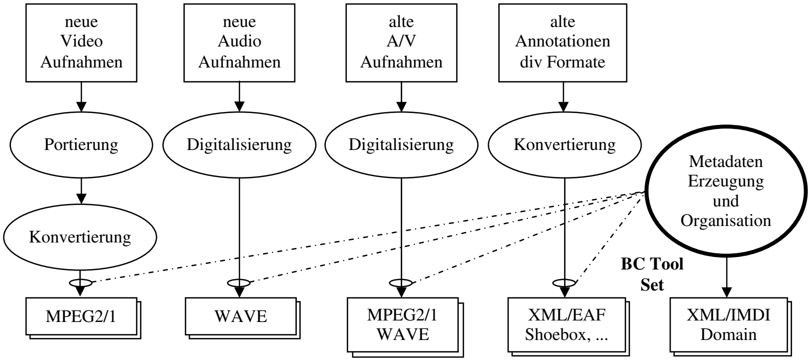

Abb. 1: Erzeugung eines organisierten Multimedialen Bestandes

In Abb. 2 wird dargestellt, daß ein Zugriff auf die Elemente des Archivs über den Metadaten Browser erfolgen kann. Wenn eine Resource gefunden wurde, kann mit verschiedenen Tools, die nach einer entsprechenden Konfigurierung im Browser verfügbar sind, sofort auf den Daten gearbeitet werden. Wenn mittels des ELAN Tools (oder anderer Tools) neue Informationen hinzugefügt werden wie z.B. eine morphologische Annotation, so muß mittels des Metadaten Editors in der Metadaten Beschreibung eine entsprechende Ergänzung vorgenommen werden. Es ist also die Metadaten Beschreibung, die von der Existenz der verschiedenen zusammengehörenden Elemente weiß.

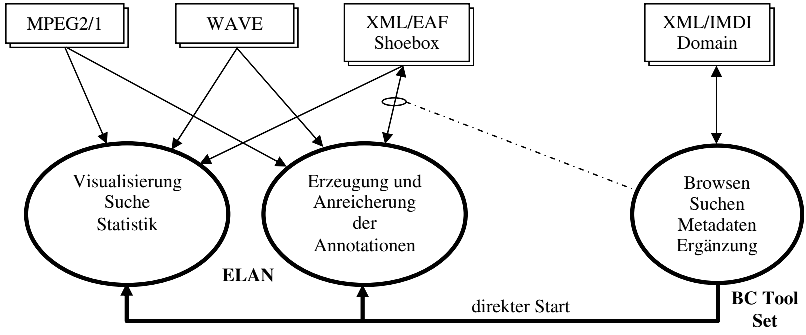

Abb. 2: Anreicherung der Bestände mit Annotationen und deren Auswertung.

4.1. Computer-basierte Audio/Video Verarbeitung

Seit ca. 3 Jahren werden am MPI alle neuen Audio und Video-Aufnahmen und auch systematisch alle älteren digitalisiert, online verfügbar gemacht und in eine Archivierungsstrategie eingebunden. Das MPI verfügt gegenwärtig über ca. 3.5 TB an Online Beständen, die in ca. 13.000 mit Metadaten-Beschreibungen versehenen Sessions (Aufnahme-Einheiten) gruppiert sind. Alle wichtigen Software-Werkzeuge, die entwickelt wurden, arbeiten sowohl im Stand-Alone Modus als auch im Netzwerk, sowohl auf normalen PCs als auch auf Notebooks. Damit ist es möglich, schrittweise Computer als alleinige Analyse Geräte zu verwenden und auf separate Audio und Video-Apparaturen weitgehenst zu verzichten. Das MPI ist auf diesem Weg bereits sehr weit vorangeschritten. Zunehmend wird Audio und Video-Technik nur noch eingesetzt, um Aufnahmen zu machen.

Sowohl für die Digitalisierung von Audio (zumeist auf DAT Recordern) als auch für Video (zumeist auf S-VHS, Hi-8 und DV Trägern) wurden eine Reihe professioneller Digitalisierungs-Setups aufgebaut, so daß unterstützt von Scripts parallel an jeweils 4 Stationen digitalisiert werden konnte. Es stellte sich schnell heraus, daß es am effizientesten war, um zunächst ein ganzes Tape zu digitalisieren und dann im Nachhinein auf dem Computer die linguistische relevanten Einheiten herauszuschneiden. Für die unterstützten digitalen Formate wurden daher Schneide- und Konversions-Programme ausgetestet und zur Verfügung gestellt. Ein genaues Workflow Schema sorgte dafür, daß die verschiedenen Files normierte Namen bekamen und in die für das Management der Daten wichtigen Dateipfade eingegliedert wurden.

4.2. Offene Standards

Von Beginn an wurde darauf orientiert, daß möglichst nur Files basierend auf Standard-Formaten ins Archiv aufgenommen werden sollten. Als Archiv-Formate wurden ausgewählt:

– XML-basierte allgemeine Formate für Metadaten, Annotationen, Lexika und andere linguistische Datentypen[19]. Für die Metadaten und Annotationen wurden mit IMDI[20] und EAF[21] Format allgemeine XML-basierende Formate entwickelt und unterstützt. Für Lexika wird gegenwärtig im europäischen Rahmen noch an einem allgemeinen Format gearbeitet. Desweiteren werden die weitverbreiteten ASCII-basierten Annotations-Formate CHAT und Shoebox[22] unterstützt. Da beide Formate eine interne Strukturierung vornehmen, können sie in das ELAN Tool eingelesen werden und auch leicht in XML umgesetzt werden.
– Für die Encodierung von Charactern wurde UNICODE gewählt. Die entwickelten Werkzeuge unterstützen UNICODE, dh. sie verfügen auch über Eingabe- und Rendering Methoden für viele Schrifttypen und -Systeme. Viele Ressourcen wurden bereits auf UNICODE übertragen, dieser Prozeß der Konvertierung bestehender Dateien wird jedoch noch andauern.
– Für die Encodierung von Video wurde MPEG2 als Archivformat gewählt, nachdem in früheren Jahren aus technischen und finanziellen Gründen mit Codecs wie MJPEG, Cinepak und MPEG1 gearbeitet wurde. MPEG1 wird vor allem aus Bandbreiten-Gründen noch stets als Frontend Format verwendet werden. In Zukunft wird für diese Zielsetzungen auch MPEG4 berücksichtigt werden.
– Für die Audio-Encodierung wird mit 44.1/48 kHz linear PCM gearbeitet, das in WAVE[23] formatierte Files abgelegt wird. Bei WAVE handelt es sich um keinen offenen Standard, allerdings wird es de facto von allen Programmen unterstützt. Komprimierte Formate wie ATRAC (Minidisc) und MP3 (MPEG Audio), die beide auf der Basis psychoakustischer Modellierungen arbeiten, werden in linear PCM expandiert. Tests auch im MPI haben ergeben, daß Informationen herausgefiltert werden, die zwar bei Sprache für das menschliche Ohr

[19] Bezüglich anderer weniger wichtiger Datentypen ist der Übergang zu XML nicht aktuell.

[20] ISLE Metadata Initiative; http://www.mpi.nl/ISLE

[21] H. Brugman, P. Wittenburg, “The application of annotation models for the construction of databases and tools – an overview and analysis of the MPI work since 1994”, Proceedings of the IRCS Workshop on Linguistic Databases, LDC, Philadelphia, December

[22] SHOEBOX (http://www.sil.org/computing/shoebox/) bietet auch ein gleichwertiges Format für Lexika.

[23] Bei WAVE handelt es sich um einen umfassenden Standard, der auch Kompressionsformate beinhaltet. Allerdings hat sich lediglich das Format für linear PCM Daten durchgesetzt.

kaum wahrnehmbar sind, jedoch bei einer späteren Bearbeitung mittels Digitaler Siganlverarbeitungstechniken zu fehlerhaften Darstellungen führen könnten.

4.3. Konvertoren

Um dem Ziel nach einem Archiv, das nur wenige offene Formate und eine UNICODE Charakter Codierung unterstützt, nahe zu kommen, mußten Konvertoren für viele andere Formate beschafft, entwickelt und integriert werden. Die Video-Format Konvertoren können aufgrund der erforderlichen Spezialkenntnisse nur beschafft werden. Es bedarf jedoch vieler Tests, um ihre Eignung in allen Umständen zu überprüfen. Für die vielen Textformate gibt es wegen der Vielfalt kaum einsetzbare Konvertoren, d.h. es mußten mehrere entwickelt werden. Am weitesten geht der Word-to-XML Konverter, der erforderlich war, weil viele Wissenschaftler ihre Annotationen in WORD ausführen und dabei eine implizite Strukturierung mittels Type-Face-Merkmalen vornehmen. Um einigermaßen flexibel zu bleiben, wurde eine Struktur-Beschreibungs-Sprache entwickelt, die es dem Wissenschaftler erlaubt, die Struktur seines WORD Dokumentes anzugeben und für die verschiedenen Elemente auch gleichzeitig die XML-Tags zu spezifizieren.

Zusätzlich wird noch eine Konvertierung vom jeweiligen Charakter-Set auf UNICODE UTF-8 vorgenommen, soweit für den Eingabe-Satz eine Abbildungsvorschrift bekannt ist.

4.4. Archiv

Im MPI wurde eine Archiv-Struktur realisiert, die die verschiedenen Datentypen trennt und je nach ihrer Art behandelt. Das Archiv basiert auf einem Hierarchical Storage Management System, das eine mittlere Tape-Library (bis zu 25 TB), einen Cache-Disk-Array umfaßt und von der SAM-FS Software organisiert wird. Für alle Dateien im Archiv werden sofort zwei Kopien angelegt. Dabei wird von den Metadaten und Text-Dateien eine Kopie immer im Online Cache gehalten. Lediglich die zwei Kopien der Media-Dateien werden ausgelagert und nur dann in den Cache geholt, wenn sie benötogt werden. Da Metadaten prinzipiell offen sind und andere Dateien jeweils mit bestimmten Zugriffsrechten versehen werden, mußte ein Dateisystem ausgearbeitet werden, welches trotz all dieser Details mit wenig Aufwand zu pflegen ist.

Da das Archiv unter anderem Aufnahmen von Sprachen und Kulturen enthält, die vom Aussterben bedroht sind und damit einen unwiderbringliches kulturellen Schatz beinhalten, muß ein Verfahren realisiert werden, das eine dritte Kopie außerhalb des Instituts-Gebäudes umfaßt. Da nur eine verteilte Speicherung, die überregionale Archivzentren einschließt, ein Überleben derartiger Daten garantiert, wird an entsprechenden Strategien gearbeitet. Diese lassen sich jedoch nur längerfristig umsetzen.

4.5. Metadaten Infrastruktur

Einer der zwei wesentlichen Stützpfeiler der Arbeit an NILE war die Entwicklung einer vollständigen Metadaten Infrastruktur, die modernen Ansprüchen und auch den Anforderungen der Linguisten genügt. Es wurde bereits erwähnt, daß der Dublin Core Metadata Set nicht ausreichend ist, um die Management Aufgaben zu lösen und um interessante Ressourcen schnell zu finden. Daher wurde aufbauend auf eine breite Übersicht[24], in Zusammenarbeit mit mehreren Feldforschern und Computerlinguisten im Rahmen des europäischen ISLE Projektes zunächst ein Metadaten Standard (IMDI) für linguistische Resourcen inklusive der Vokabulare für die verschiedenen Elemente entwickelt. Dieser IMDI Standard wurde inzwischen auf der Basis zweier Workshops mit internationalen Experten zur Version 3.2 weiterentwickelt. Er enthält gegenwärtig Vorschläge zur Beschreibung von Korpora, Lexika und geschriebenen Ressourcen.

Um komplexe Fragen beantworten zu können, wurde IMDI als ein strukturierter Metadaten-Satz entworfen. Teilblöcke, die sich oft wiederholen wie z.B. Projektbeschreibungen, können separat gespeichert und wieder eingebunden werden. Sehr wichtig für das Management von und den Zugriff auf Ressourcen ist das Konzept der Bündelung. Es sind die Metadaten die die enge Kopplung diverser Ressourcen, wie z.B. die in Abb. 3 dargestellten, realisieren.

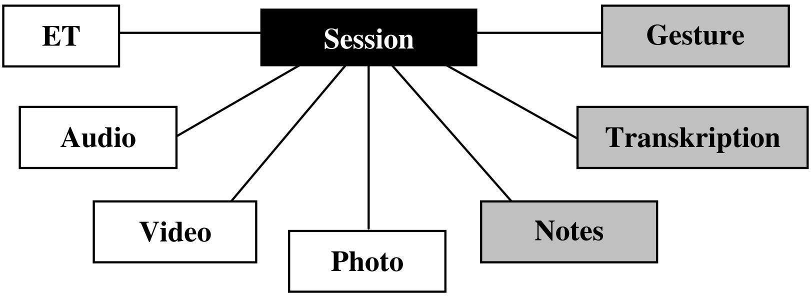

Abb. 3 zeigt eine Reihe von eng zusammengehörender Dateien. Die linken Blöcke zeigen die zu einer bestimmten Aufnahmen gehörenden Rohdaten, die parallel aufgenommen wurden, während die dunklen Blöcke die Annotationen darstellen. Die Metadaten enthalten die Informationen über deren Zusammengehörigkeit.

[24] http://www.mpi.nl/ISLE/

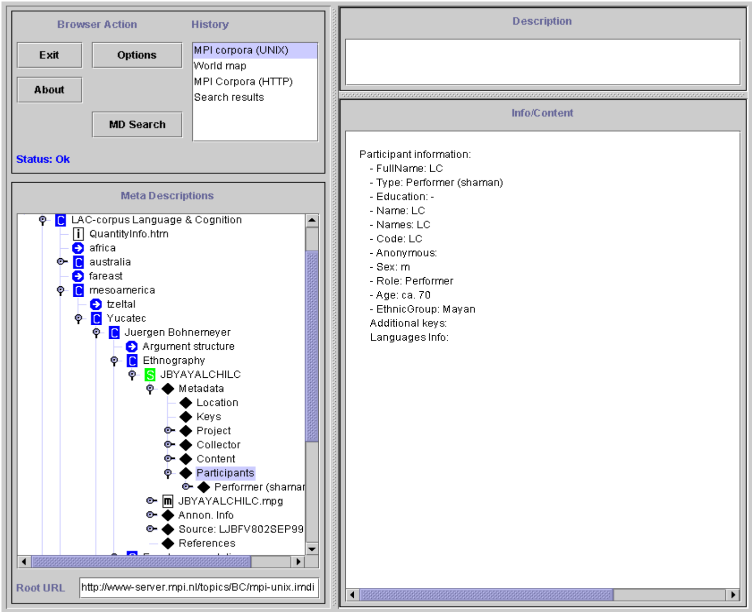

Abb. 4 zeigt den Metadaten Browser. Im Bild unten links ist ein Teil einer konkreten Baumstruktur widergegeben, so wie sie von den betreffenden Wissenschaftlern definiert wurde. Die grünen Icons zeigen an, dass eine Session gefunden wurde, die durch Metadaten (schwarze Diamant Symbole) beschrieben ist und die auch einen MPEG File (mpg) enthält. Im großen rechten Fenster können verschiedene Informationen angezeigt werden, die zu dem betreffenden Knoten gehören. In diesem Falle sind einige Metadaten Elemente angezeigt.

Im Gegensatz zu vielen anderen Initiativen war von vornherein beabsichtigt, die IMDI Domäne so zu gestalten, daß sie nicht nur das Suchen, sondern auch das Browsen in Hierarchien (s. Abb. 4) erlaubt. Zum Korpusmanagement benötigt man eine kanonische Struktur, allerdings sollte es jedem Benutzter erlaubt sein, Suchresultate oder Unterbäume wiederum zu seiner neuen temporären Domäne zu machen. Das Vorhandensein derartiger Bäume erleichtert das Eingliedern verschiedener Korpora, die durch verschiedene Institute bzw. sogar durch Individuen zur Verfügung gestellt werden. Lediglich die Top-URLs müssen registriert werden, um eine distribuierte Domäne zu erzeugen und es ist auch für Computer-Laien einfach, um den eigenen Bestand zu integrieren. Die Metadaten Bäume dienen ebenfalls dazu, beliebige Informationsdateien mit Knoten verknüpfen. Eine der Hierarchie-Ebenen in einem solchen Baum wird sicherlich durch die behandelten Sprache definiert. So macht es Sinn, Beschreibungen der Grammatik und des Lautsystems einer Sprache mit den entsprechenden Knoten zu verbinden. Auch aufgrund eines einfachen Managements von Dateien ist eine derartige Verlinkung sinnvoll.

Ein professioneller Editor wurde entwickelt, der den vollen IMDI Standard, ein Caching und das Verwenden der Vokabulare, die Wiederverwendung von Teildaten und eine Reihe von nützlichen Funktionen wie das Anzeigen der Definitionen unterstützt. Er erzeugt validierte XML Daten. Ein Browser erlaubt das Navigieren in dieser Domäne verlinkter und distribuierter Metadaten-Beschreibungen. Er analysiert die XML-Inhalte und kann diese wie auch gelinkte Informations Dateien im HTML oder PDF Format präsentieren. Er erlaubt es ebenfalls, eine Konfigurations-Datei mit Programm-Namen so anzulegen, daß die auf einem bestimmten File-Typ möglichen Operationen angezeigt werden und direkt ausgeführt werden können. Es obliegt dem Benutzer, die entsprechenden bevorzugten Programme einzutragen. Wie bereits erwähnt wurde, kann der Benutzer Teilbäume oder Suchresultate als neue Domäne markieren oder zusammenstellen, um sich so für eine gewisse Zeit eine effiziente Arbeitsumgebung zu erzeugen.

Eine Suchkomponente wurde programmiert und in den Browser integriert, die es erlaubt, in derartigen verteilten Räumen nach Metadaten Mustern zu suchen. Aus Effizienzgründen wird erwartet, daß mittels einfacher Kommandos pro registrierter IMDI Domäne eine Suchdatei angelegt wird. Da das Parsen aller XML-Dateien zu Zeitaufwendig wäre, schickt das Suchprogramm die Suchfragen an die entsprechenden IMDI Instanzen, die ihrerseits lediglich eine Suche in der indexierten Suchdatei ausführt. Bisher wurde auf ein zentrales Harvesting und das Abspeichern aller Metadaten in einer speziellen XML- oder relationalen DB verzichtet. Diese könnte durch die neuen Projekte mit größer werdenden Metadaten-Repositorien jedoch unabdingbar werden, um die Suchresultate schnell präsentieren zu können.

Um die IMDI Metadaten Beschreibungen Service-Providern gegenüber zu öffnen, deren Services auf Dublin Core basieren, wurde ein Mapping zwischen IMDI und Dublin Core vorgenommen[25] und das Harvesting-Protocol der Open Archives Initiative[26] implementiert. Bei der Übertragung können viele IMDI Details wegen der eingeschränkten Semantik von DC nur üngenügend durchgegeben werden, d.h. die DC Benutzer erhalten nur eine bschränkte Sicht.

Darüberhinaus wurden verschiedene Scripte erstellt, um IMDI Files z.B. aus Spreadsheet-artigen Beschreibungen zu erzeugen und z.B. sehr viele Metadaten-Beschreibungen mit einem einzigen Befehl anzupassen. Solange diese Scripts jedoch nicht in die Browser und Editor Umgebung integriert sind und dadurch eine Validitäts-Kontrolle bezüglich der erzeugten Daten vorgenommen wird, können diese Tools nicht an die Benutzer weitergegeben werden und bleiben dem Korpus-Management vorbehalten.

[25] http://www.mpi.nl/ISLE
[26] Open Archives Initiative; http://www.openarchives.org/

Im gerade gestarteten INTERA Projekt[27]

- wird eine Europa-weite IMDI Domäne erzeugt, indem sich verschiedene Zentren mit Sprachressourcen zusammenschließen werden,
- werden alle Elemente des IMDI Metadaten Satzes entsprechend geltender ISO Normen definiert und diese Definition werden in offenen Repositories verfügbar sein,
- werden alle impliziten Relationen mittels RDF[28] explizit gemacht, so daß Teilblöcke und Elemente von anderen wiederverwendet und beliebige Services von anderen erzeugt werden können,
- wird die IMDI Ressource Domäne derart mit Repositories verbunden, in denen Tools beschrieben sind, daß es möglich gemacht werden soll, daß Benutzer in Abhängigkeit der gefundenen Ressource Typen beliebige im Tool Repository beschriebene Tools starten können.

Insbesondere letzteres wird mittels moderner Web-Services Verfahren zusammen mit den Partnern realisiert werden.

Im ebenfalls gerade gestarteten ECHO[29] Projekt wird das Metadaten Konzept in Richtung auf Multidisziplinarität und Interoperabilität erweitert werden müssen, d.h. semantischen Aspekten wird eine bedeutendere Rolle zukommen.

Innerhalb des im Mai 2002 gegründeten ISO Kommitees ISO TC37/SC4 wird eine Untergruppe Standards für Metadaten-Beschreibungen für linguistische Ressourcen ausarbeiten. Die Erfahrungen mit IMDI werden darin Berücksichtigung finden.

4.6. ELAN Annotations und Analyse Tool

Der andere wesentliche Stützpfeiler der Arbeiten am MPI war die Entwicklung eines Tool Sets zum Arbeiten mit multimedialen/multimodalen Ressourcen auf Rechnern, um somit Audio/Video Apparatur nicht nur überflüssig zu machen, sondern darüberhinaus eine integrierte Annotations- und Analyse Umgebung für Files zu bieten, die es erlaubt, verschiedene Medien wie Audio, Video, Augenbewegungssignale und anderes mehr gleichwertig zu behandeln und mit Texten zu verknüpfen. Auf der Basis einiger Werkzeuge der ersten Generation wie z.B. MediaTagger wurde das EUDICO Tool Set Konzept entwickelt und schrittweise implementiert. Inzwischen wurden das ELAN Annotations Tool, verschiedene Viewer und ein Suchtool entwickelt.

[27] Integrated European (Language Resource Area; http://www.mpi.nl/INTERA
[28] Resource Description Framework; http://www.w3.org/RDF/
[29] European Cultural Heritage Online; http://www.mpi.nl/echo

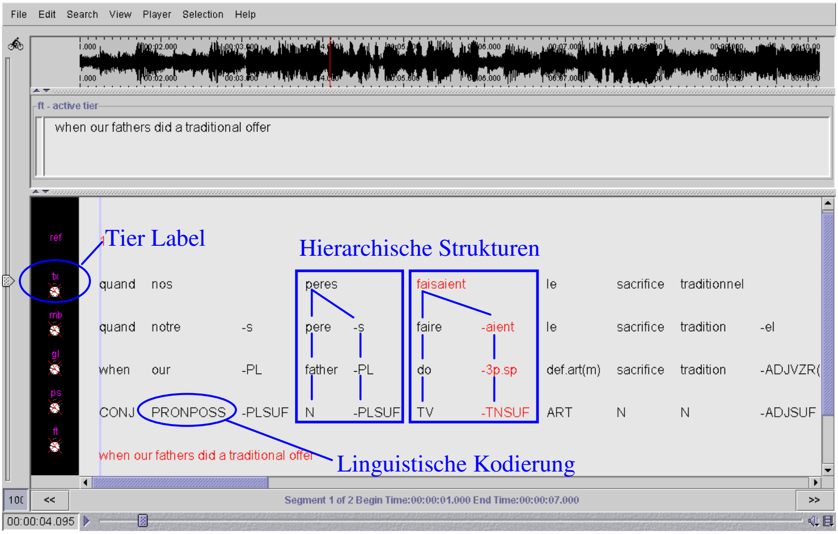

Abb. 5 zeigt eine typische Darstellung in ELAN, wobei ein Sprachfile mit hierarchischen Annotationen versehen worden ist. Die groß eingetragenen Elemente sind zur Erklärung hinzugefügt. Ganz links im schwarz unterlegten Panel sind die 6 vom Benutzer frei gewählten Tier Label zu sehen. Im großen Fenster sind die einzelnen Annotationen in den verschiedenen Schichten zu erkennen. Angedeutet sind zwei Beispiele hierarchischer Kodierungen.

Letztere wurden schrittweise in ELAN integriert, um eine vollständige und daher für den Benutzer einfach handhabbare Umgebung anzubieten.

Basis für die Entwicklung des EUDICO Tool Sets war die Ausarbeitung und Fortschreibung des Abstract Corpus Model (ACM)[30]. Dadurch, daß es ein allgemeines Model für einen multimodalen Korpus inklusive der komplexen Annotationen beschreibt, definiert es den Kern der Klassenstruktur der EUDICO Tools. Ebenfalls gestattet es, verschiedene bekannte Formate wie z.B. Shoebox oder CHAT (Annotations Format in CHILDES project) einzulesen und auf das ACM abzubilden. Die Verfügbarkeit derartiger Import-Module garantiert die format-unabhängige Arbeit der Wissenschaftler. Mit dem EAF (EUDICO Annotation Format) wurde ein XML-basiertes flexibles Annotations-Format erzeugt, das es erlaubt, die während der Arbeit mit ELAN erzeugten Daten persistent zu machen. Dieses Format bietet die Möglichkeit, Benutzer-definierte Tiertypen zu speichern, Abhängigkeiten zwischen Tiers darzustellen, Annotationen flexibel an die Zeit zu

[30] H. Brugman, P. Wittenburg, “The application of annotation models for the construction of databases and tools - an overview and analysis of the MPI work since 1994”, Proceedings of the IRCS Workshop on Linguistic Databases, LDC, Philadelphia, December

koppeln, Annotationen aber auch flexibel an mehrere bestehende Annotationen zu linken und Hierarchien darzustellen.

Das ELAN Annotations-Tool, das jetzt in der Version 1.3 verfügbar ist und aufgrund von Benutzerwünschen in seiner Handhabung optimiert werden konnte, bietet eine ganze Reihe interessanter und innovativer Funktionen wie z.B.

- Benutzer-definierbare Tiers inkl. deren Abhängigkeiten,
- hierarchische Annotationen wie sie z.B. für Morphologie erforderlich sind (siehe Bild 4),
- Annotationen in verschiedenen Character-Sets wie z.B. Chinesisch, Arabisch, Hebräisch, IPA inkl. Support für deren Schriftsysteme (Ligaturen beim Arabischen),
- leichtes Markieren eines Medien-Fragmentes, dessen auditiver bzw. visueller Kontrolle und dessen Annotierung,
- Veränderung der Präsentationsgeschwindigkeit bei Audio und Video Signalen,
- verschiedene Viewer für die Annotationen, wobei alle Viewer derart verknüpft sind, daß ein Navigieren in einem Viewer auch zu einem Markieren der entsprechenden Elemente in den anderen führt,
- Unterstützung zum schnellen Einlesen und zur Zeitzuordnung bereits existierender Annotationen.

Für Analyse-Aufgaben steht ein Suchtool zur Verfügung, das es ermöglicht

- nach komplexen Mustern zu suchen, wobei mehrere mit bestimmten Tiers assoziierte Einzel-Muster logisch miteinander verknüpft werden können,
- eine Verknüpfung ist auch über zeitliche und textuelle Abstände möglich, wobei die definierten Abhängigkeiten und Hierarchien berücksichtigt werden,
- die gefundenen Resultate direkt zum Navigieren zu verwenden, dh. es kann sofort auf die entsprechende Stelle im Text gesprungen werden.

Weitere Möglichkeiten sind im Handbuch[31] beschrieben.

Gegenwärtig wird im Rahmen zweier Projekte an einer Ergänzung des Codes in zwei Richtungen gearbeitet: (1) Es wird in Zukunft eine Selbstverständlichkeit für Benutzer sein, um gleichzeitig an einem Korpus zu arbeiten, obwohl sie an verschiedenen Stellen arbeiten. Das Tool muß also kollaboratives Arbeiten unterstützen, ohne die Konsistenz der Daten zu gefährden. (2) Es muß um eine Lexikon-Schnittstelle ergänzt werden, die eine Interaktion zwischen Lexikon und Korpus ermöglicht.

Das EUDICO Tool Set ist momentan im Bereich der multimedia Language Resources sicherlich eines der am weistesten entwickelten Werkzeuge.

---

[31] http://www.mpi.nl/tools

4.7. Tools Site

Alle Tools sind auf der Tools Web-Seite des MPI [www.mpi.nl/tools] frei verfügbar. Mittels des Webstart Tools (JNLP)[32] können alle vom MPI fertiggestellten Tools in einfacher Weise über das Web heruntergeladen und auch gestartet werden. Auf dieser Web-Seite sind ebenfalls alle Handbücher verfügbar, die im Detail über die Funktionalität Auskunft geben. Weiterhin ist ein multimediales Beispiel-File vorhanden, das bereits mit multimodalen Annotationen (Sprache und Gestik) versehen ist, so daß das ELAN Tool getestet werden kann.

Alle Tools werden auf der W2000 Plattform programmiert. Wegen der verwendeten Java Version sind die Browsable Corpus Tools (Editor und Browser) nur auf MAC OS10 lauffähig. Das ELAN Tool wird gegenwärtig derart umgeschrieben, daß neben Java Media Framework auch Quicktime[33] unterstützt wird. Diese Arbeiten wird in Kürze abgeschlossen sein, so daß ELAN dann unter MAC OS10 lauffähig sein wird.

Die Browsable Corpus Tools sind auch unter Linux ausführbar. Bezüglich ELAN muß jedoch festgestellt werden, daß es noch keinen adäquaten Media Player gibt.

Es ist eine Bug-Report Datenbank programmiert worden, so daß Beutzer ihre Reports eingeben und auch den Stand der Behandlung erfragen können. Diese wurde mittels einer Datenbank realisiert, die über eine Web-Form ansprechbar ist.

4.8. Korpus

Auf den mittels der beschriebenen Methoden erzeugten Korpus kann über die Web-Seite „www.mpi.nl/corpus“ zugegriffen werden. Die Metadaten Beschreibungen aller Ressourcen sind offen zugänglich, d.h. mittels des IMDI Browsers und bei Einstellung des richtigen Bookmarks kann der Bestand betrachtet werden. Der Zugang auf die Ressourcen selbst ist aus ethischen und legalen Gründen zumeist gesperrt und kann nur über wissenschaftliche Kooperationen und entsprechende formale Vereinbarungen erreicht werden. Die IMDI Domäne umfaßt auch Korpora von verschiedenen Projekten wie z.B. dem DOBES Projekt, dem ESF Zweitspracherwerbs Projekt und dem Dutch Spoken Corpus Projekt.

[32] Webstart; http://java.sun.com/products/javawebstart/
[33] Quicktime; http:// www.apple.com/quicktime

5. Ausblick

Innerhalb eines relativ kurzen Zeitrahmens wurde eine vollständige Infrastruktur zum Aufbau, zur Analyse und zur Verwaltung eines Digitalen Multimedialen Korpus aufgebaut, die den Anforderungen der Mitarbeiter des Max-Planck-Institutes, aber darüber hinaus auch denen anderer wichtiger Projekte und individueller Wissenschaftler entspricht. Im Rahmen europäischer Projekte wurden auch Grundpfeiler implementiert, die eine Öffnung in Richtung auf das Semantische Web ermöglichen und Aspekte eines modernen, kollaboralen Arbeitens beinhalten. Innerhalb neuer Projekte werden diese Ideen weiter ausgebaut.

Das INTERA Projekt (Integrated European Language Resource Area) umfaßt die folgenden Aufgaben für das MPI Entwicklerteam: (1) Es muß eine Metadaten Domäne aufgebaut werden, die eine ganze Reihe wesentlicher europäischer Datenzentren und zentrale Portale und Repositorien umfaßt. (2) Es werden alle Element und Vokabular Definitionen, die zum jetzigen IMDI Metadaten Satz gehören, mittels der durch ISO 11179 und ISO 16642 definierten Verfahren explizit formuliert und in offenen Repositorien abgelegt. Ebenso werden alle impliziten Relationen mittels RDFS formuliert und alle Schemata (XML und RDF) offen abgelegt. (3) Weiterhin wird eine Verbindung zwischen einem Tool Repositorium, indem Tools mittels geeigneter Metadaten beschrieben werden, und der IMDI Ressourcen Domäne derart aufgebaut, daß ein Wissenschaftler, der geeignete Ressourcen gefunden hat, sofort ausführbare Tools angezeigt bekommt und diese dann auch ausführen kann. (4) Der Zugang auf die genannten Informationen und der Informationsaustausch zwischen Tool- und Ressource Repositorien soll mittels Web-Services erfolgen, d.h. die Services werden mittels UDDI Registraturen bekannt gemacht.

Im ECHO Projekt (European Cultural Heritage Online) soll eine Disziplin-übergreifende Metadaten Such-Domäne etabliert werden, was neben syntaktischen (die Daten sind in verschiedenen Formaten vorhanden) vor allem semantische Probleme aufwirft, die zu lösen sind. Darüberhinaus wird das ELAN Tool derart erweitert, daß ein kollaboratives Arbeiten über das Internet ermöglicht wird, d.h. daß der bereits vorhandene Kerncode um entsprechende durchaus nicht-triviale Mechanismen erweitert werden muß.

Alle diese Arbeiten stellen weitere Schritte zu mehr Flexibilität und in Richtung des Semantischen Web dar. So erlaubt die explizite Darstellung von Relationen und Definitionen entsprechend anerkannter Standards es anderen, eigene Services bzw. Applikationen zu realisieren, die Teile der angebotenen Informationen mit denen anderer Repositorien verknüpfen. Durch diese Projekte wird gewährleistet, daß die bisher entwickelten Methoden und Tools im Sinne neuer Trends weiterentwickelt und damit auch weiterhin gepflegt werden.

Nominiert für den Heinz-Billing-Preis 2002

gene: A Tool for Studying Turbulent Transport in Fusion Plasmas

Frank Jenko
Max-Planck-Institut für Plasmaphysik, Garching

Abstract

The gene code (this acronym stands for **g**yrokinetic **e**lectromagnetic **n**umerical **e**xperiment) is a tool to help tackle the "Grand Challenge" problem of plasma microturbulence from a fully kinetic point of view. To this aim, the nonlinear gyrokinetic equations for a magnetized electron-ion plasma are solved on a fixed grid in five-dimensional phase space, employing finite-difference and pseudo-spectral techniques. All relevant electromagnetic effects are taken into account. Each simulation requires of the order of 10^9 grid points and up to 10^5 time steps, challenging even the most powerful present-day computers. gene runs efficiently on multiple massively parallel platforms, achieving, e.g., 280 MFlops on the Hitachi SR-8000 at Leibniz Computing Center at Munich. To reach this level of performance, it was necessary to adapt the implementation strategy to each individual machine architecture. With the help of gene simulations, it has recently been demonstrated that turbulence in magnetized fusion plasmas may in fact extend all the way down to the smallest possible spatial scales. By analyzing a large number of state-of-the-art simulations, it could be explained why this fascinating phenomenon had been overlooked in the past. These results opened up a new and presently very active area of computational plasma physics.

1 The Challenge

Research in the area of magnetic confinement fusion aims at exploring the physical processes underlying the development of a novel energy source. Using the sun as a role model, the idea is to heat up a gas consisting of deuterium

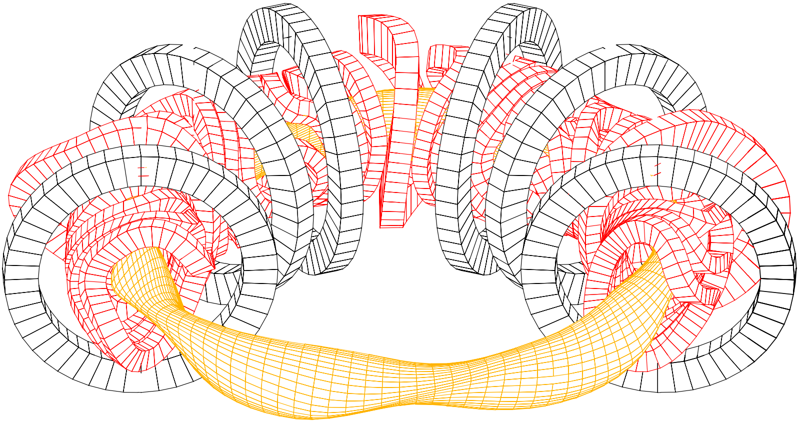

Fig. 1: Schematic view of the fusion experiment Wendelstein 7-AS at the Max Planck Institute for Plasma Physics (IPP) at Garching. Shown are some of the external magnetic field coils and a representative "magnetic surface" spanned by a single magnetic field line.

and tritium to some 100 million degrees. At this enormous temperature, the probability of energy releasing fusion reactions between these two different hydogen isotopes is sufficient to envision a fusion power plant. Of course, no material boundary could withstand such conditions, and therefore this hot ionized gas - or plasma - is instead confined by a toroidal (i.e., doughnut-shaped) magnetic field (see Fig. 1). Since the charged plasma particles basically follow the magnetic field lines, they can thus be kept away from the wall. For sufficiently large plasma pressure and energy confinement time, the rate of fusion reactions is large enough to simultaneously heat the plasma and allow for energy to be extracted. However, small-scale ("microscopic") plasma waves can become unstable in the presence of density and temperature gradients, leading to plasma turbulence associated with strongly enhanced particle and heat transport. This imperfect magnetic confinement can only be compensated for by building bigger and thus more expensive devices. It is therefore one of the main obstacles on the way to a future fusion reactor.

Turbulent flows are ubiquitous in nature and science. From leaves dancing in the autumn wind to the dynamics of the interstellar medium, turbulence is a fascinating and often important topic to be studied. Moreover, its understanding is key to the development of better aircrafts or automobiles. From a more fundamental point of view, turbulence is a classic example of nonlinear dynamics in open systems. Here, the system typically establishes a quasistationary state far from thermodynamic equilibrium. For this to happen, a permanent input, redistribution, and output of energy is required. Turbulence is often described as a bath of vortices of varying sizes and lifetimes, sometimes spanning several orders of magnitude, and all of them nonlinearly created, coupled, and destroyed. In view of these facts, it might not come as a surprise that the Nobel Prize Laureate Richard Feynman called turbulence "the most important unsolved problem of classical physics."

Several generations of physicists have attempted - with moderate success - to tackle this problem purely analytically, but with the recent advent of powerful supercomputers, turbulent flows became accessible to new and promising forms of analysis. Numerical experiments yield insights into its structure and dynamics, and sometimes they even allow for quantitative predictions. However, plasma turbulence is even more complex and computationally demanding than ordinary fluid turbulence. E.g., whereas the latter is described by the so-called Navier-Stokes equations of fluid dynamics, magnetically confined fusion plasmas are generally characterized by very long mean free paths, thus violating some of the assumptions underlying standard fluid theory. This means that first-principles simulations need to be based on a gyrokinetic treatment [1] of multiple particle species in five-dimensional phase space. (The fast gyromotion around magnetic field lines can be removed from the basic equations analytically, reducing the number of dimensions from six to five. Hence, the label "gyrokinetic".) Moreover, due to a zoo of linear and nonlinear drive mechanisms, plasma turbulence involves active perpendicular (throughout this text, "perpendicular" and "parallel" are defined with respect to the confining magnetic field) spatial scales ranging from the ion gyroradius of a few mm all the way to the system size of several m. At least, that was the conventional wisdom until a few years ago.

2 Surprising Scientific Results

Since 1999, both experiments and simulation results from the newly developed `gene` code indicated that the turbulent activity may indeed extend all the way down to the so-called Debye length. For typical experimental conditions, the latter is of the order of the electron gyroradius, i.e., less than 0.1 mm. It represents the smallest length scale in a plasma which is able to sustain microinstabilities and therefore turbulence. The possible relevance of these "hyperfine" scale fluctuations came as a big surprise, mainly because widely used theories based on Ludwig Prandtl's mixing length concept predict that their contribution to the total turbulent transport is negligible. In the light of our numerical results, this view was questioned, however. Since then, the nonlinear physics responsible for this completely unexpected finding has been investigated and some of its key ingredients have been disclosed. It is crucial to note that this progress could only be made through the effective use of massively parallel computers. But before we present the software tool that has been designed and employed for this purpose, we would like to briefly review some key physical results and insights in the remainder of this section. The reader interested in more details might want to consult one or more of the following publications: [2, 3, 4, 5, 6, 7, 8, 9].

2.1 Temperature gradient driven microinstabilities

A magnetized plasma is subject to many microscopic instabilities. A prominent example is the temperature gradient driven (TG) mode. It exists in two versions, the ITG mode and the ETG mode, driven, respectively, by gradients in the ion and electron temperature. The dynamics of TG modes can generally be described as perpendicular destabilization competing against parallel delocalization. Since a rigorous treatment of TG modes is fairly involved, we will make use of simple analogies. To understand the perpendicular dynamics, we allude to the following situation: An ideal fluid is stratified in terms of its mass density ρ such that ρ decreases with height. This system exhibits neutrally stable waves with frequency $\omega = (g/L_\rho)^{1/2}$ where g is the gravitational constant and $L_\rho = |d \ln \rho / dx|^{-1}$ is the mass density scale length. If ρ *increases* with height, we obtain an exponentially growing instability with growth rate $\gamma = (g/L_\rho)^{1/2}$. This observation can be readily transferred to fusion physics. Plasma particles confined in a toroidal magnetic field are subject to a centrifugal force which can be written in terms of an effective gravitational constant as $g_{\text{eff}} \sim v_t^2/R$ where v_t is the thermal velocity of the particles and R is the major radius of the toroidal device. On the outboard side, the effective gravity points in the direction of decreasing plasma pressure, yielding an "interchange" instability with the growth rate $\gamma \sim (g_{\text{eff}}/L_T)^{1/2} \sim v_t/(RL_T)^{1/2}$. Of course, the particles move along the helically wounded field lines and thus average over different values of the growth rate (which vanishes on the inboard side). Therefore a TG mode depends on mode localization on the outboard side which must be provided by the parallel dynamics. This situation is similar to how one can prevent honey from dripping from a honey dipper by turning it sufficiently fast: There is an instability if and only if the parallel transit frequency is smaller than the interchange growth rate, $\gamma \gtrsim v_t/R$ or $R/L_T \gtrsim 1$, a condition confirmed by more quantitative analyses.

2.2 Are ITG and ETG turbulence isomorphic?

Basic ITG and ETG modes are perfectly isomorphic, i.e., linear results obtained for one of these microinstabilities can also be applied to the other one by simply interchanging the species labels.[2] E.g., the intrinsic perpendicular length scale is given by the ion and electron gyroradius, respectively. The

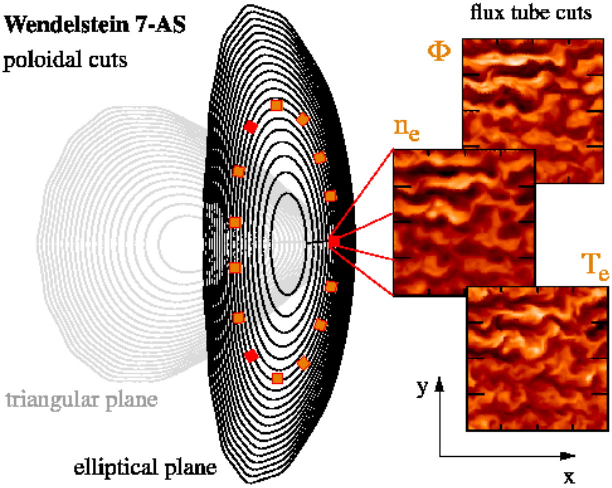

Fig. 2: Cross sections of the fusion experiment Wendelstein 7-AS at the Max Planck Institute for Plasma Physics (IPP) at Garching. The simulations are done in representative flux tubes (following the helically wounded magnetic field lines) and exhibit radially elongated streamers.

latter is smaller than the former by the square-root of the ion-to-electron mass ratio, a number that is typically in the range of 40-60 for fusion experiments. In the nonlinear regime, this symmetry is broken, however. To understand why, it is helpful to introduce the concept of secondary instabilities. A linear instability grows exponentially until the gradients of its spatially varying mode amplitudes are so large that it in turn triggers a secondary instability. At this point, the mode nonlinearly saturates. The transition to turbulence has occurred. It can be shown that for TG modes, there are two distinct types of secondaries, one of which differs for ITG and ETG modes.[8, 9] Consequently, under certain conditions the ETG simulation goes to much higher normalized saturation levels than its ITG counterpart and exhibits radially elongated vortices called “streamers” (see Fig. 2). This represents an interesting example of structure formation in turbulent flows and raises the associated transport way beyond mixing length expectations. In fact, with this enhancement, ETG-induced transport can be comparable to electron thermal transport induced by ITG modes on much larger spatial scales. This result is completely unexpected and strongly counterintuitive.

2.3 On the nature of ETG streamers

The single most striking feature in nonlinear ETG simulations is thus the occurrence of radially elongated vortices with large fluctuation amplitudes. These streamers have been observed for different magnetic configurations, but only in parameter regimes which favor such structures even linearly and are not subject to a strong secondary instability.[2, 3, 6, 7] For large aspect ratio tokamaks with circular, concentric magnetic surfaces, this is the case for a magnetic shear value (quantifying the radial variation of the winding number of magnetic field lines) of $\hat{s} \gtrsim 0.4$ and $R/L_{T_e} \gg R/L_{T_e}^{\mathrm{crit}}$.[2] In the presence of streamers, the fluctuation and transport levels can be boosted by more than an order of magnitude with respect to mixing length expectations.[8] Characteristic statistical properties of the fully developed turbulence like correlation functions, streamer aspect ratios, and phase shift relations have also been determined.[2, 6, 7] The poloidal wavenumber spectra exhibit a power law behavior with exponents that seem to be quite universal.[7, 8] The induced turbulent transport (both with *and* without streamers) is clearly dominated by its electrostatic part (associated with perpendicular $\mathbf{E} \times \mathbf{B}$ advection), with the electromagnetic component (due to fast electron motion along perturbed magnetic field lines) contributing only some 1% to the total electron heat diffusivity.[2, 7] This finding is again in contrast to expectations created by oversimplified analytic theories since the 1970's.

2.4 On the role of turbulence-generated shear flows

Because of the strong external magnetic field, turbulence in magnetized plasmas is very anisotropic (typical parallel to perpendicular correlation length ratios are of the order of 10^3-10^5) and sometimes resembles two-dimensional fluid turbulence in perpendicular planes. Under these conditions, the turbulent energy cascade is directed towards *larger* spatial scales describing nonlinear structure formation through *merging* vortices. As is known, e.g., from the study of planetary atmospheres, this effect may drive zonal shear flows. In magnetoplasmas, zonal flows are linearly stable but can be self-generated by the turbulence and may in turn act as its dominant nonlinear saturation mechanism. But how important are they on electron gyroradius scales? It is known that low-amplitude ITG streamers can be broken up by zonal flows.[10] Here, the radial shearing rate of the zonal $\mathbf{E} \times \mathbf{B}$ flow is found to exceed the maximum linear growth rate of ITG modes, leading to vortex decorrelation and turbulence suppression. Since there is currently no comprehensive analytic theory predicting the behavior of zonal flows, numerical simulations have to reveal their role and nature in the ETG regime.[7] They are found to be 15-20 times weaker than in the ITG case, failing to affect the high-amplitude streamers. Moreover, zonal fields – the magnetic analogon of zonal flows – are also too small to play a significant role in the turbulent dynamics. Once again, the spectral properties of zonal modes are observed to be quite universal, indicating a close connection to cascade physics.

2.5 Semi-analytical modeling of turbulent transport by ETG streamers

Thus, first-principles nonlinear computations have led to a series of discoveries which have overthrown previous intuitive expectations and analytic predictions. Building on these new results and insights, an attempt can be made to capture some of the key features of the turbulent dynamics in an *a posteriori* semi-analytical model. In streamer dominated regimes, i.e., for $\hat{s} \gtrsim 0.4$ and $R/L_{T_e} \gg R/L_{T_e}^{\rm crit}$, the growth rates of ETG modes and associated secondaries can in fact be condensed into relatively simple algebraic formulas.[9] Since the secondary growth rates are proportional to the amplitude of the underlying ETG modes, assuming a balance between the two growth rates in fully developed turbulence determines the fluctuation (and transport) levels,

$$\chi_e \approx \mathcal{F}(\cdots)\,\frac{R}{L_{T_e}}\,\frac{\rho_e^2 v_{te}}{L_{T_e}} \tag{1}$$

where $\mathcal{F}$ is a function of a few local plasma parameters. The application of secondary instability theory of course presupposes that the back reaction of zonal flows and fields on the turbulence can be ignored. But as we have stated above, this is generally the case on ETG scales. We find that streamer-dominated turbulent transport is reasonably well described by this model (given the enormous complexity of the problem) as can be seen in Fig. 3. Importantly, several conventional estimates for turbulent transport even fail to capture the observed *trends*. Therefore, this can be considered as a nice example of the successful interaction between simulation and theory. The question if this approximative approach can be generalized to non-streamer ETG regimes or to other forms of plasma turbulence is currently under investigation.

2.6 The new frontier: Cross-scale coupling

For many years, the focus of plasma turbulence research has been on length and time scales set by the heavier ions. The study of ITG modes falls into this category of “fine” scale turbulence. ETG modes, on the other hand, are characterized by “hyperfine” scales which are smaller by a factor of the square root of the ion-to-electron mass ratio, i.e., by a factor of about 40 for hydrogenic plasmas. Having demonstrated that despite these disparate scales,

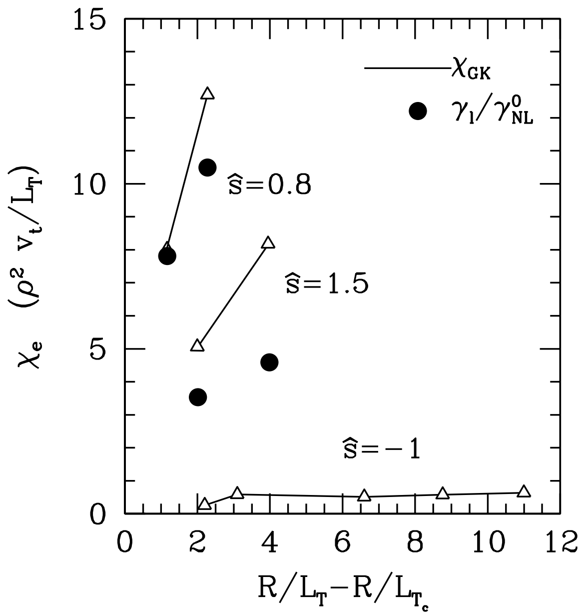

Fig. 3: Electron heat flux from nonlinear gyrokinetic simulations of ETG turbulence together with semi-analytical predictions. In streamer dominated regimes, one finds reasonable agreement.

ETG-induced transport can compete with ion-scale transport [2], the question of cross-scale coupling naturally arises. E.g., due to fine-scale turbulence, the electron temperature gradient seen by ETG modes can be substantially different from the value set by the equilibrium profiles. On the other hand, the presence of hyperfine-scale turbulence may even effect the linear properties of ITG modes. Many other nonlinear coupling scenarios can be envisioned.

It is probably needless to say that the simulations required here represent an extremely demanding computational task, pushing the employed software and hardware to their limits. Nevertheless, such studies have great potential for the discovery of new physical phenomena as indicated by the first numerical results for cross-scale coupled ETG-ITG turbulence shown in Fig. 4 (see also Ref. [7]). A long-term, synergistic effort is likely to be needed to deeply penetrate this exciting area of plasma turbulence research. We expect that `gene` will continue to serve as a useful tool in the context of this ambitious enterprise.

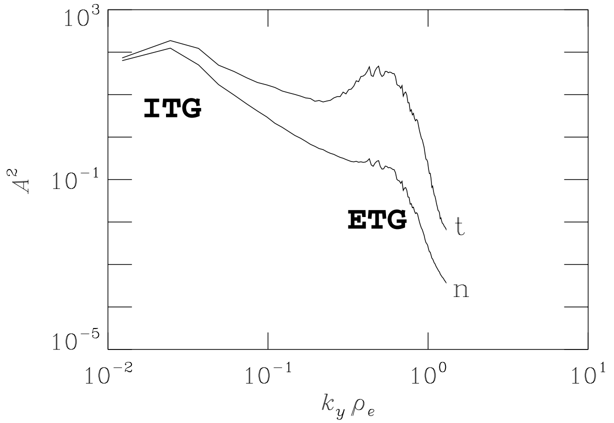

Fig. 4: Perpendicular wave number spectra of temperature and density fluctuations from first gyrokinetic simulations of cross-scale coupled ETG-ITG turbulence. Novel nonlinear effects are expected to be encountered in this system.

3 The Tool

The studies reviewed in the last section only became possible with the help the `gene` code which allows to perform well-resolved nonlinear simulations of ETG and cross-scale coupled turbulence on large space-time domains. Although other gyrokinetic codes do exist, due to its rather efficient phase-space grid algorithm and very good parallel performance, `gene` is uniquely positioned to address such demanding problems of computational plasma turbulence research. Moreover, it is presently one of only two gyrokinetic codes that is able to go beyond the commonly used adiabatic electron approximation and treat magnetic field field fluctuations and, in particular, zonal fields. Obviously, efficiency is key to make phenomena like the ones described above numerically accessible. This requires optimizations on various levels, including the choice of the algorithm and its practical implementation on massively parallel computers. These aspects shall be briefly highlighted in this section.

3.1 Algorithm

Algorithms for solving kinetic equations can be roughly devided into two groups according to the Lagrangian and Eulerian description of phase space dynamics:

(a) Particle codes, which are widely used in computational plasma physics, sample the phase space dynamics by following the trajectories of an ensemble of N_p “superparticles”. However, the resulting numerical noise, scaling like $N_p^{-1/2}$, can be reduced but not eliminated by extracting part of the dynamics analytically (see Ref. [11] and references therein). For this reason, the total run time of turbulence simulations tends to be severely restricted by noise build-up, leading to the violation of fundamental conservation laws. Beyond that, the inclusion of magnetic field fluctuations has also proven to be difficult. In order to capture the important dynamics of kinetic shear Alfvén waves, two terms in the modified gyrokinetic Ampère's law have to cancel exactly. The inherent noise inhibits this cancellation unless the spatial resolution is taken to numerically prohibitive limits. Recent attempts to circumvent this problem [12] have yet to be implemented and tested under realistic conditions. Moreover, the gyroaverging of the electromagnetic fields is done by an N-point technique [13] which is not suited for cross-scale simulations. We therefore conclude that particle techniques do not seem to be the ideal choice for the problems under consideration.

(b) A complementary approach is to adopt a Eulerian point of view and represent the gyrocenter distribution functions on a fixed grid in five-dimensional phase space. The nonlinear gyrokinetic equations [1] are then finite differenced appropriately. Since in many relevant situations (including streamer-dominated ETG turbulence), the time step in a plasma turbulence simulation is set by the nonlinearities, anyway, one does not gain anything by using time-implicit methods (see, e.g., the discussion in Ref. [11]). We therefore opt for a time-explicit scheme based on multidimensional second-order upwind methods [14] which is less complex, more efficient, and easier to parallelize (for more details, see Ref. [2] and references therein). The time step is adjusted dynamically. By means of a special parallel canonical momentum technique, this noiseless scheme has no problems in reproducing the dynamics of kinetic shear Alfvén waves. In addition, the effect of fast particles (important for computing higher-order velocity space moments) and strong local variations in velocity space (say, around Landau resonances or the trapped-passing particle boundary) can be captured well by adaptive techniques. Despite their reputation of being computationally expensive, grid-based gyrokinetic turbulence codes can often compete with their particle counterparts in terms of the net effort to solve a given problem.[11]

3.2 Massively parallel implementation

A typical ETG turbulence simulation requires of the order of 10^9 grid points and up to 10^5 time steps. This is to ensure the system has reached a quasi-stationary nonequilibrium state involving all relevant degrees of freedom and has sustained it long enough to yield good statistics. (Note that compared to other kinds of plasma turbulence, ETG simulations are characterized by larger radial box sizes and smaller time steps to account for the presence of

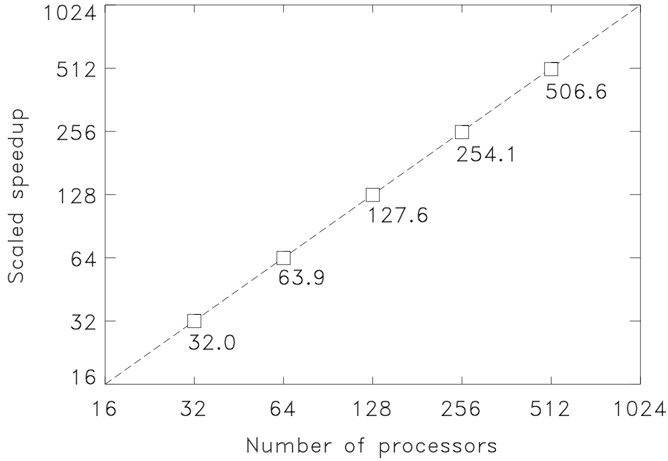

Fig. 5: Parallel performance of an earlier version of `gene` *on the Cray T3E-600 at the Computing Center Garching. For comparison, the ideal speedup curve is shown as a dashed line.*

high-amplitude streamers.) Hence, a massively parallel implementation is clearly called for. E.g., the run shown in Fig. 3 of Ref. [5] took 150 h on 256 processors of the Hitachi SR-8000 and used about 0.5 GByte of memory per PE. Typical output files are in the GByte range. These numbers will increase further as various aspects of the cross-scale coupling problem are addressed.

`gene` runs on multiple platforms, including the Cray T3E-600 and a new IBM SMP (Regatta) system at Rechenzentrum Garching as well as the Hitachi SR-8000 at Leibniz-Rechenzentrum München. It has been found that in order to make effective use of these machines, it is absolutely necessary to adapt the parallelization strategy to each individual architecture. On the Cray, the SHMEM library is used in concert with spatial domain decomposition techniques to obtain some 60 MFlops. Here, a speedup of 506 on 512 PEs has been measured for a slightly modified (so-called "drift-kinetic") version of `gene` (see Fig. 5).[11] For the Hitachi, `gene`'s communication routines were rewritten using MPI, and parts of the code had to be customized for optimal single PE performance. Thus, 280 MFlops have been achieved. However, porting the code to the IBM, which consists of multiple SMP nodes, a fundamental change in the parallelization strategy became inevitable. In the new version of the code, four out of five dimensions (including all three spatial dimensions) are kept on a node, while one of the two velocity space dimensions is distributed across the nodes. Thus the cross-node communication is greatly reduced, leaving only a few global sums per time step. To this aim, a combination of MPI and OpenMP has been used. Single PE performance

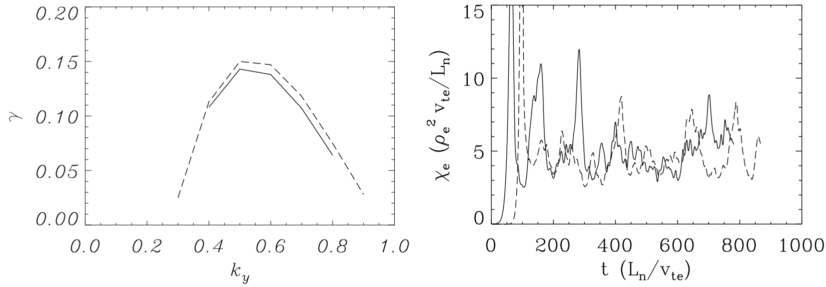

Fig. 6: Comparison of linear growth rates of ETG modes and time traces of turbulent transport (for a non-streamer ETG case) from `gene` *(solid lines) and* `gs2` *(dashed lines).*

was again optimized. In all versions of the code, memory usage was cut down to an irreducible minimum. Moreover, the time to write and read restart files has been kept reasonable levels (of the order of a few minutes) through parallel I/O. Despite all of the above options, `gene` is still maintained as a single code, using a small number of precompiler switches to invoke them.

3.3 Benchmark tests

For a code as involved as `gene`, it is vital to benchmark it against other codes before attempting to utilize it in large-scale studies. Such tests have been performed both in the linear and in the nonlinear regime. Examples can be seen in Fig. 6, in which linear growth rates of ETG modes and time traces of turbulent transport (for a non-streamer ETG case) from `gene` and another gyrokinetic code, `gs2` [15], are shown for comparison. Note that since turbulence is a chaotic phenomenon, one can only expect nonlinear results to agree in statistical sense. This is the case here.

4 Summary

The possible role of hyperfine-scale turbulent transport in fusion plasmas has been explored through a large number of state-of-the-art simulations with the software tool `gene`. In the course of this project, a series of unexpected and fascinating discoveries were made, opening up a new and presently very active area of computational plasma turbulence research. To make these advances possible, `gene` was optimized to run at 10-20% of the theoretical peak performance on multiple massively parallel platforms. With the full installation of the new IBM SMP (Regatta) system at Rechenzentrum Garching, `gene` will be utilized to help pioneer the emerging research area of cross-scale coupled plasma turbulence, pushing the employed software and hardware to their limits.

Acknowledgements

I would like to thank my colleague W. Dorland for fruitful scientific collaborations, and R. Tisma at the Computing Center Garching for help with performance tuning.

References

[1] E. A. Frieman and L. Chen, Phys. Fluids **25**, 502 (1982).
[2] F. Jenko, W. Dorland, M. Kotschenreuther, and B. N. Rogers, Phys. Plasmas **7**, 1904 (2000).
[3] W. Dorland, F. Jenko, M. Kotschenreuther, and B. N. Rogers, Phys. Rev. Lett. **85**, 5579 (2000).
[4] F. Jenko, W. Dorland, and G. W. Hammett, Phys. Plasmas **8**, 4096 (2001).
[5] F. Jenko and W. Dorland, Plasma Phys. Control. Fusion **43**, A141 (2001).
[6] F. Jenko and A. Kendl, New Journal of Physics **4**, 35 (2002).
[7] F. Jenko and A. Kendl, Physics of Plasmas **9**, 4103 (2002).
[8] F. Jenko and W. Dorland, Phys. Rev. Lett. **89**, 225001 (2002).
[9] F. Jenko, W. Dorland, B. Scott, and D. Strintzi, in *Theory of Fusion Plasmas 2002*, p. 157.
[10] Z. Lin, T. S. Hahm, W. W. Lee *et al.*, Science **281**, 1835 (1998).
[11] F. Jenko, Comput. Phys. Commun. **125**, 196 (2000).
[12] Z. Lin and L. Chen, Phys. Plasmas **8**, 1447 (2001).
[13] W. W. Lee, J. Comput. Phys. **72**, 243 (1987).
[14] P. Colella, J. Comput. Phys. **87**, 171 (1990).
[15] M. Kotschenreuther, G. Rewoldt, and W. M. Tang, Comput. Phys. Commun. **88**, 128 (1995).

Introducing `LambdaTensor1.0` – A package for explicit symbolic and numeric Lie algebra and Lie group calculations

Thomas Fischbacher
Max-Planck-Institut für Gravitationsphysik,
Albert-Einstein-Institut, Potsdam

Abstract

Due to the occurrence of large exceptional Lie groups in supergravity, calculations involving explicit Lie algebra and Lie group element manipulations easily become very complicated and hence also error-prone if done by hand. Research on the extremal structure of maximal gauged supergravity theories in various dimensions sparked the development of a library for efficient abstract multilinear algebra calculations involving sparse and non-sparse higher-rank tensors, which is presented here.

1 Introduction

Although nowadays, there are many powerful general-purpose symbolic algebra packages available, the most well-known ones probably being `Mathematica`, `Maple`, `MuPAD`, `MACSYMA`, and `Reduce`, there still are many open problems which require (sometimes highly specialized) computer algebra that has become feasible on modern machines, but is beyond the scope and limitations of these systems. Hence, in some areas of research, there is notable demand for tailored symbolic algebra.

In this work, a package for group-theoretic calculations is presented which originally was created for the computation of explicit expressions for the potentials of gauged maximal supergravity theories on Euler-angle-like parametrized submanifolds of the manifolds of supergravity scalars. Currently, this seems to be the most effective tool for the determination of these potentials and hence possible vacua of these supergravity theories, but it may be applicable to a wider class of more general group-theoretic problems. Nevertheless, the structure of mentioned supergravity potentials turned out to be so rich that they have not been mined completely up to date, which by itself is a strong incentive to make this package publicly available.

There are two complementary ways how such a package could be implemented: either as a library manipulating tensor expressions on the symbolic level, like `MathTensor`, or as a library operating on explicit tensor coordinates (which nevertheless may well be symbolic themselves). Here, the latter approach was chosen for two reasons: first, the possibility to always go down to plain numbers on the fly leads to better testability and broader applicability in the field of numerical calculations, and second, since one typical application is to break the exceptional group $E_{8(+8)}$ to subgroups as small as possible, the number of terms generated on the purely symbolic level would explode, nullifying the conceptual advantages of the former approach.

Some of the outstanding features of the symbolic algebra framework presented here are:

– It is Free Software, available under the GNU Lesser General Public License, Version 2.1. (Nevertheless, users are urged to read the accompanying copyright information to see how it is applied in this particular case.)
– It is complementary to the `LiE` package in that sense that `LiE` can not perform calculations that deal with explicit Lie algebra or Lie group elements, while such functionality lies at the heart of `LambdaTensor1.0`.
– It is reasonably efficient, since it employs and is designed for an optimizing machine code compiler.
– It is extensible *at the level of the implementation itself*, not, like most other symbolic algebra packages, only at the level of a built-in application language. In particular, *it does not restrict the user of this package to a simple stripped-down C-like-in-spirit application language*. Instead, the user has full access to the feature-rich, extensible, compiler system on which this package was built, CMU Common LISP.

It is perhaps noteworthy that utilizing a functional subsystem (in this case, Common LISP) not only obsoletes the need for the implementation of yet another application-centric programming language, which hardly could and should compete with a system backed by its own user community and team of developers, but also readily makes available a variety of other useful existing libraries and tools written for this system. The rationale behind the design decision to use Common LISP, and in particular to build on the CMU CL implementation, instead of using one of the other viable alternatives suggesting themselves here, namely Haskell (in particular, GHC), Objective Caml, or Scheme (in particular, Gambit), which is of course mainly a question of personal preferences of the author, comes in part from the experienced inflexibility of systems based on typed λ-calculus (ML and Haskell – although the power of Haskell's typeclass system has to be acknowledged), partly due to general typability problems, partly since these systems do not provide means to modify the very language itself which would be comparable to the macro facilities of LISP (although this is much less a problem in a lazy functional language like Haskell), and in part from the baroque richness in features of the Common LISP standard (especially when compared to Scheme). The main drawback of this decision is that the most advanced freely available Common LISP compiler system, CMU CL, is only well supported for x86-based platforms, and due to technical reasons, comes with a limitation on the size of the data being processed in memory of 800 MB[1].

The central idea behind this package is to utilize the observation that tensors showing up in group theory calculations frequently are very sparsely occupied – for example, using the conventions of [2], structure constants of the largest (248-dimensional) exceptional Lie group E_8 $f_{\mathcal{AB}}{}^{\mathcal{C}}$ contain only 49 440 out of $248^3 = 15\,252\,992$ nonzero entries – and hence, we can make good use of efficient implementations of abstract algorithms that can handle sparsely occupied higher-rank tensors. Efficient code working on sparse matrices is widely used and readily available; the appropriate algorithms for handling higher-rank tensors are also quite well-known, albeit in a very different context: relational databases.

In particular, at the level of explicit tensor entries, forming a quantity like

$$M_{abc} = N_{gha}P^{gh}{}_{bc} \tag{1.1}$$

translates as follows into the language of relational databases (SQL syntax used here):

```
SELECT t1.index3 as index1,
       t2.index3 as index2,
       t2.index4 as index3,
       SUM(t1.val*t2.val)
   FROM tensor1 t1, tensor2 t2
   WHERE t1.index1=t2.index1 AND t1.index2=t2.index2
   GROUP BY t1.index3, t2.index3, t2.index4;
```

Unfortunately, it is not feasible to just connect to an existing SQL database system (like PostgreSQL), create relations for tensors, and use existing imple-

[1]Experienced Unix users can raise this to about 1.6 GB, but in many cases this requires in particular manual modification of the kernel source.

mentations of these algorithms by doing all the calculations in the database, for various reasons. Besides considerations concerning the efficiency of communication, and considerable additional computational overhead due to databases having different aims, one major problem is that extending the database system to abstract from the implementation of sum and product here, as is necessary as soon as we want to work with data types not natively supported by the database (which are frequently limited to integers and floatingpoint numbers) would bring along too many technical problems. Hence, what is required is a re-implementation of the underlying database algorithms with numerical and symbolic tensor computations as applications in mind. Furthermore, this implementation has to be abstract enough to allow all relevant arithmetic operations to be provided as parameters, so that one may switch between approximate numerics, exact (i.e. rational number) numerics, and symbolic calculations. (The ability to implement and use arbitrary arithmetics on tensor entries has proven to be of great value during the debugging phase of the symbolic algebras provided within this package. For example, it is easy to lift an existing implementation of arithmetic operations on symbolic terms to an implementation working on pairs of terms and numerical values of these terms for a given occupation of variables that signals an error whenever a discrepancy between these values shows up.)

The `LambdaTensor1.0` package consists of the following parts:

1. Various general-purpose functions providing important infrastructure. (containing e.g. simple combinatorial functions, basic linear algebra, balanced binary trees, simple optimization functions, priority queues, and a serializer.) For the Debian GNU/Linux system, this part is available as a separate package.
2. Support for sparse higher-rank arrays and tensor operations on them, where implementations of arithmetic operations on tensor entries may be given as parameters.
3. Different implementations of symbolic algebra. (One which is similar in spirit and intention, though not in scope, to conventional general-purpose symbolic algebra packages, one which is aggressively optimized (and hence far from being general-purpose) for calculations involving products of trigonometric functions of the particular form showing up in supergravity calculations as in [1, 2], and a third one utilizing the CMUCL port of the MAXIMA symbolic algebra package.)
4. Applications. In particular, definitions relevant for the exceptional groups $E_{7(+7)}$ and $E_{8(+8)}$ and the potentials of maximal gauged extended supergravity theories in three and four dimensions.
5. Worked out, documented examples that demonstrate how to use the package.

Since this package was created as a byproduct of work targeted at the determination of the extremal structure of supergravity potentials, these tools are in some aspects just as good as they had to be for this task, with lots of opportunities for optimization and improvement still remaining. In a different vein, since those particular calculations are quite demanding, these tools *are* quite optimized in the most central aspects. Nevertheless, large parts of this codebase are constantly exchanged, improved, re-written, and hence, major changes should be expected between version 1.0 and subsequent versions.

2 An overview over `LambdaTensor1.0`

Since detailed technical documentation can be found in the manual accompanying the library, we only want to give a brief overview over concepts and algorithms underlying the different pieces of this package.

2.1 General purpose functions

This is a collection of various functions and macros ranging from simple more convenient redefinitions of features already available in COMMON LISP to implementations of ubiquitous algorithms to complex facilities providing vital infrastructure for the other parts of `LambdaTensor1.0`. Since parts of this highly inhomogeneous conglomerate of functions and definitions are not essential for `LambdaTensor1.0`, but perhaps useful in a much broader context, the decision was made to split this off into a separate package for the the Debian distribution of `LambdaTensor1.0`. Current functionality provided here encompasses, but is not limited to, macros providing machine-code optimization information to the compiler, various abbreviations and definitions that were inspired by the Perl language, elementary combinatorial functions, basic polynome factorization, linear algebra, and optimization support, as well as efficient implementations of balanced binary trees and priority queues.

2.2 Sparse array functions

This is the heart of `LambdaTensor1.0`, implementing sparse arbitrary-rank arrays using database algorithms. In particular, a sparse array is represented internally as a multidimensional hash of its nonzero components. Sparse arrays are transparently re-hashed if their occupation density grows, up to a certain percentage, where the implementation internally switches to storing tensor entries in a nonsparse array. Currently, removing entries from a sparse-array does not induce the underlying hash to shrink once occupation density falls below a certain level.

The most important sparse array function provided is `SP-X` which implements efficient tensor multiplication, contraction, and index reordering of an arbitrary number of tensors (limited by resources) where multiplications and contractions are heuristically sequenced in such a way to minimize the total number of operations. To give an example of what `SP-X` can do and how it is used, let's assume that the variable `F-abc` contains the structure constants ${f_{ab}}^c$ of a Lie algebra. Then, the Cartan-Killing metric $g_{ab} = {f_{ap}}^q {f_{bq}}^p$ can be computed as follows:

```
(defvar metric
   (sp-x '(a b)
         '(,F-abc a p q)
         '(,F-abc b q p)))
```

A second example: assuming `so8-sigma` is the rank-3 tensor of $SO(8)$ Γ-matrices with index order $(i, \alpha, \dot\beta)$, the following piece of code checks the Clifford algebra properties under contraction of the cospinor indices:

```
(sp-multiple-p
 (sp+ (sp-x '(i j a b)
            '(,so8-sigma i a a*)
            '(,so8-sigma j b a*))
      (sp-x '(i j a b)
            '(,so8-sigma j a a*)
            '(,so8-sigma i b a*)))
 (sp-x '(i j a b)
       '(,(sp-id 8) i j)
       '(,(sp-id 8) a b)))
```

Besides tensor arithmetics (parametrized by the underlying implementation of arithmetics on tensor entries)[2] and functions computing embedding tensors for index split operations, this piece of code also provides a selection of linear algebra functions on sparse tensors (like `SP-MATRIX-EIGENVALUES` and `SP-INVERT` for quadratic rank-2 tensors over numbers (not yet symbolic expressions)), conversion functions mapping sparse arrays to and from non-sparse vectors, as well as a collection of group theory related functions like `SP-FIND-ROOT-OPERATOR` which, given structure constants, a Cartan subalgebra, and a root vector, determines the adjoint-representation coefficients of the corresponding ladder operator, or `SP-LIN-INDEP-COMMUTATORS`, giving a linearly independent basis for all the commutators of two sets of quadratic sparse rank-2 tensors.

[2] for some kinds of tensor entries, in particular complex double-precision floatingpoint numbers, the implementation uses specially optimized versions of internal functions.

2.3 Symbolic Algebra

There are different implementations of symbolic algebra available within this package, each of them having its own raison d'être. The most effective since most highly optimized towards the problem for the original task of computation of supergravity potentials is the *packof-exp*, or, in brief, *po-exp* algebra. When looking at a typical supergravity potential restricted to a gauge subgroup singlet manifold, like the following one from [2] of $N = 16, D = 3$ supergravity with gauge group $SO(8) \times SO(8)$ on the conveniently parametrized four-dimensional $(SL(2)/U(1))^2$ manifold of $G_{2,\mathrm{diag}}$ singlets,

$$\begin{aligned} -8g^{-2}V \quad = \quad & \tfrac{243}{8} + \tfrac{7}{2}\cosh(2\,s) + \tfrac{49}{8}\cosh(4\,s) + \tfrac{1141}{64}\cosh(s)\cosh(z) \\ & + \tfrac{427}{64}\cosh(3\,s)\cosh(z) - \tfrac{7}{64}\cosh(5\,s)\cosh(z) \\ & - \tfrac{25}{64}\cosh(7\,s)\cosh(z) + \tfrac{21}{8}\cos(4\,v) \\ & - \tfrac{7}{2}\cos(4\,v)\cosh(2\,s) + \tfrac{7}{8}\cos(4\,v)\cosh(4\,s) \\ & - \tfrac{21}{64}\cos(4\,v)\cosh(s)\cosh(z) \\ & + \tfrac{21}{64}\cos(4\,v)\cosh(3\,s)\cosh(z) \\ & + \tfrac{7}{64}\cos(4\,v)\cosh(5\,s)\cosh(z) \\ & - \tfrac{7}{64}\cos(4\,v)\cosh(7\,s)\cosh(z) \\ & - \tfrac{1645}{128}\cos(v - w)\sinh(z)\sinh(s) \\ & + \tfrac{651}{128}\cos(v - w)\sinh(z)\sinh(3\,s) \\ & + \tfrac{7}{128}\cos(v - w)\sinh(z)\sinh(5\,s) \\ & - \tfrac{49}{128}\cos(v - w)\sinh(z)\sinh(7\,s) \\ & - \tfrac{315}{64}\cos(3\,v + w)\sinh(z)\sinh(s) \\ & + \tfrac{133}{64}\cos(3\,v + w)\sinh(z)\sinh(3\,s) \\ & - \tfrac{7}{64}\cos(3\,v + w)\sinh(z)\sinh(5\,s) \\ & - \tfrac{7}{64}\cos(3\,v + w)\sinh(z)\sinh(7\,s) \\ & + \tfrac{35}{128}\cos(7\,v + w)\sinh(z)\sinh(s) \\ & - \tfrac{21}{128}\cos(7\,v + w)\sinh(z)\sinh(3\,s) \\ & + \tfrac{7}{128}\cos(7\,v + w)\sinh(z)\sinh(5\,s) \\ & - \tfrac{1}{128}\cos(7\,v + w)\sinh(z)\sinh(7\,s), \end{aligned} \tag{2.1}$$

one notices that the typical summand in such a term (as well as in all intermediate quantities) is of the form $k \cdot \exp\left(\sum_j c_j v_j\right)$, where v_j are variable names, and c_j as well as k all are either real or imaginary rational numbers. Furthermore, these terms 'come in packs' and can be grouped together to form summands like $\frac{7}{128}\cos(7\,v + w)\sinh(z)\sinh(5\,s)$. Hence, we can forge an internal representation of such terms which is orders of magnitude more efficient both in terms of memory consumption and possible reductions than the conventional one of a generic term (as used by other computer algebra systems) by using this additional structure. (Details are given in the manual.)

It is the combination of this problem-specific implementation of a symbolic algebra with efficient sparse array database algorithms that made it possible to transcend all previous limitations in complexity in [2].

As is perhaps imaginable, this aggressively optimized symbolic algebra was not the first one to be employed in conjunction with sparse tensor algorithms. The former one, which is available as the function-polynome (*funpoly*, or often briefly *fp*) algebra, is much closer in design to conventional symbolic algebra, hence also more flexible and still shows up in some places within `LambdaTensor1.0`. (Note that the *poexp* algebra is so specialized that it can not handle anything else but terms of the structure described above. In particular, it can not represent quasipolynomials, hence one easily runs into trouble when nilpotence enters the stage.) The funpoly symbolic algebra also implements some peculiarities going beyond what one may expect from conventional symbolic algebra packages, in particular some non-local reductions of precisely that kind which FORM avoids by construction in order to be able to efficiently handle formulae much bigger than available memory. Since this piece of code did not undergo as many evolutionary cycles of being rewritten as some of the rest of this package, its design still shows some flaws and weaknesses[3], and so it is scheduled for replacement in later versions.

Since it may nevertheless be important to also have a flexible, general, tested, and powerful implementation of some symbolic algebra available that can be used in conjunction with this package, even if one should not try to use that particular one for the calculation of supergravity potentials, this library also comes with a simple interface to the free `MACSYMA`-replacement `MAXIMA`, which was originally implemented on top of `GCL`, but then also ported to other LISPs, including CMU CL.

2.4 Applications and examples

`LambdaTensor1.0` comes with optional additional definitions related to the groups $E_{8(8)}$, $E_{7(7)}$ as well as important subgroups thereof and further functions relevant for the computation and investigation of the structure of the scalar potentials of three- and four-dimensional maximal gauged supergravity theories. Finally, detailed worked-out examples are provided within the package which explain how to apply it to supergravity calculations.

3 Availability and concluding remarks

The most recent version of `LambdaTensor` is available from the webpage
`http://www.cip.physik.uni-muenchen.de/~tf/lambdatensor/`

[3]in particular, handling of fractional powers of rational numbers is quite clumsy

To the author's best knowledge, the library presented here is the first abstract implementation of efficient fundamental sparse higher-rank tensor multilinear algebra, thus possibly closing an important gap. Therefore, the author considers a release under a free software license, in particular, the GNU LGPL 2.1, as adequate. As stated in detail in the accompanying copyright information, users of this library are asked to quote the present article, since it is customary to use citations as a rough measure for scientific relevance, so that further development of this work can be kept up to the needs of its users.

4 Addendum: New features in version 1.1

Since an update of this software package hardly justifies a new paper, a brief overview over new functionality introduced in version 1.1, released on 25.03.2003, is given in this addendum. First and perhaps most important, the installation process has been considerably simplified, now providing packages for all major Linux distributions. Besides some minor bugfixes to code and documentation, a considerable amount of new functionality has been implemented to extend the group-theoretic capabilities of this library into the direction of the LiE program. Among the major new algorithm implementations are a version of the Fast Freudenthal algorithm to calculate weight multiplicities of representations of simple groups, the Peterson recursion formula to determine root multiplicities of Kac-Moody algebras (this was used in [3] to calculate level decompositions of the infinite-dimensional algebras E_{10} and E_{11}), as well as LISP-oriented versions of the Todd-Coxeter coset enumeration algorithm and the Schreier-Sims algorithm for permutation groups.

References

[1] T. Fischbacher, Nucl. Phys. B **638** (2002) 207 [arXiv:hep-th/0201030].

[2] T. Fischbacher, H. Nicolai and H. Samtleben, Class. Quant. Grav. **19** (2002) 5297 [arXiv:hep-th/0207206].

[3] H. Nicolai and T. Fischbacher, arXiv:hep-th/0301017.

Weitere Beiträge für den Heinz-Billing-Preis 2002

The Virtual Laboratory Infrastructure for Controlled Online Experiments in Economics

Ben Greiner*, Hans–Arno Jacobsen†, Carsten Schmidt‡

*Max Planck Institute for Research into Economic Systems, Strategic Interaction Group, Jena
†Department of Electrical and Computer Engineering and Department of Computer Science, University of Toronto, Canada
‡Corresponding author: Max Planck Institute for Research into Economic Systems, Strategic Interaction Group, Jena

Abstract

The goal of this paper is to provide an overview on the *Virtual Laboratory* infrastructure for controlled online experiments in economics. We summarize our experience gained from performing several economic experiments on the Internet. The experiments we have run range from electronic markets to individual decision making. From there we synthesize and evaluate the methodological issue of experimental control in performing economic experiments on the Internet. The paper discusses IT-based solutions to maximize control over subjects and the environment when conducting experiments online. As a result for further exploration we sketch the design of an infrastructure that allows the automated execution of Internet experiments including marketing of experiments, control of application and participation, payment system integration, and evaluation of results.

Keywords: Internet experiments, Internet services, experimental economics, methodology.

1 Introduction

The Internet provides a natural testbed to conduct experiments with human subjects. Early psychological online studies proved that it was not only possible to conduct research online (Krantz, Ballard & Scher 1997), but that it was also feasible to collect large samples of quality data in a short period of time (Birnbaum 2000). Inspired by these early advances in psychological online research experimental economists started to run experiments via the Internet (Budimir & Rieck 1998, Baier, Bolle, Buschbaum & Swiniarska 1997, Lucking-Reiley 1999, Anderhub, Müller & Schmidt 2001, Shavit, Sonsino & Benzion 2001, Drehmann, Oechssler & Roider 2002).

Internet experiments have various advantages over computer–based laboratory experiments:

1. higher participation rates
2. feasibility to conduct experiments with long duration (e.g., days)
3. access to a more diverse subject pool (demographically, culturally)
4. higher ecological validity (artificial laboratory vs. familiar environment)
5. avoid experimenter effects
6. feasibility to conduct experiments without an expensive laboratory setup
7. automation of many experimenter tasks
8. Internet experiments run in the laboratory, but laboratory experiments do not run via the Internet, at least most of the time

The main drawback of the open environment of the Internet is that the experimenter looses some control. One aim of this paper is to provide IT-based solutions to some of the following problems

1. less control of subjects (double participation, group decision, drop-out)
2. less control of environment (use of aids, quality of network connection)
3. immediate payment (currently not feasible)

Internet experiments can be conducted in a more or less controlled environment. Some experiments are conducted close to the standards of laboratory experiments (Anderhub et al. 2001, Shavit et al. 2001) with values induced to participants (Smith 1976) and computerized user interface, some are conducted less controlled with home grown preferences and e-mail communication (Lucking-Reiley 1999). It seems that many variations between both extremes are possible. We are aware that the proposed IT-solutions to increase control might exclude participants and might not attract a representative subject pool. Most experimental economists are more concerned about control than about representativity of the subject pool since the theories which are tested in general do not make predictions with regards to the demographic distribution of the economic agents (Camerer 1997).[1] Still, the subject pools

[1] Online research in psychology is more concerned about representativity of their subjects.

of Internet experiments are far more diverse than the usual student subjects in the lab (see Anderhub et al. (2001) for details).

The goal of this paper is to provide a technical guide for experimental economists who want to conduct a controlled online experiment. This paper is not about on how to implement an online experiment, other authors have dealt with this topic (Kirchkamp 2000), rather on how to technically design and organize an economic Internet experiment in the light of maximizing experimental control.

Other online infrastructures for economic experiments exist. The Veconlab at the University of Virginia provides a service to run economic experiments for interactive learning. The site offers about 30 different experiments which can be parameterized by the teacher in order to run an economic experiment for teaching purposes (Holt 2002).[2] Further, the Iowa electronic markets[3] (Forsythe, Nelson, Neumann & Wright 1992, Forsythe, Rietz & Ross 1999) and the AuctionBot[4] (Wurman, Wellmann & Walsh 1998) provide an infrastructure for a very specialized set of economic experiments, i.e., trading based on political and economic events and auctions, respectively. An early attempt to provide a more general experiment infrastructure for research purposes is the Vlab at the University of California, Berkeley, that seems to be currently unmaintained.[5] There are numerous psychological online experiment sites on the Internet. As a good starting point the interested reader might visit the (virtual) experimental psychology laboratory.[6]

Section 2 discusses the methodological issue of experimental control in online economic experiments. In Section 3 IT-based solutions that provide maximum control for the experimenter are presented. Finally, Section 4 sketches the design rational of the *Internet experiment infrastructure*, we have developed, and Section 5 draws conclusions.

2 Experimental control

This section discusses the methodological issue experimental control with regards to Internet experiments. The discussion relies on methodological findings of online research in experimental psychology that are reviewed in the light of experimental economics.

The control of Internet experiments can be distinguished in three types of questions (Reips 1997): (1) preventing subjects from cheating, (2) con-

[2] `http://www.people.virginia.edu/~cah2k/programs.html`

[3] `http://www.biz.uiowa.edu/iem`

[4] `http://tac.eecs.umich.edu/auction`

[5] `http://elsa.berkeley.edu/vlab`

[6] `http://www.psychologie.unizh.ch/genpsy/Ulf/Lab/WebExpPsyLab.html`

trolling variables in the sense of experimenting in a controllable laboratory, and (3) avoiding confusion. Whereas the latter question is not different from laboratory experiments, the first two questions are of main concern, both to participants and researchers. Table 1 gives an overview on the items a controlled economic experiment via the Internet can satisfy with respect to the control of subjects and the environment.

Requirements	Internet	laboratory
Prevent subjects from cheating		
A subject should not play twice	yes	yes
Decisions are not made by a group but by individuals	no	yes
A subject should not contact former subjects who did the experiment before	partly	partly
Controlling variables in the sense of experimenting in a controllable environment		
Only controllable help devices are used	partly	yes
Control of subject interaction with GUI	partly	yes
Control drop-out of participants	yes	yes
Participants should take the recommended time to solve the problem	yes	yes
Control quality of the network connection	partly	yes
Payment of subjects right after the experiment	no	yes

Tab. 1: Satisfiable requirements for controlled experiments

A major concern related to the Internet technology is to ensure that subjects do not play twice. Internet protocols and Internet services do not provide the feature of unique identification of subjects, so far. We propose the following. For a controlled economic environment the use of a reward medium, usually money, is a self evident percept (Friedman & Sunder 1994). Nearly all electronic payment systems have the built-in feature of identification of payer and payee, a feature that Internet protocols do not provide. By using the identification mechanism of the payment system to identify subjects, the problem of double participation can be solved in an elegant manner.

More difficult to control is that decisions are made by each subject individually and not by groups. In contrast to laboratory environments it is impossible to control if there is more than one person involved in the decision process. Similar to the laboratory environment the problem persists that subjects contact former subjects who did the experiment before.[7] Here, a short

[7]The Internet provides several communication channels, like newsgroups and chat, which

period of time the experiment is available on the Internet provides some help to prevent subject communication. In addition, the recruitment of a heterogeneous subject pool might be advantageous.

Controlling the environment of the subject when conducting an online experiment is our second major concern. The experimenter cannot control whether aids were used to solve the task. Information technology provides the option to offer additional help devices to the subjects, for example a calculator, that can be controlled by the experimenter. A problem related to the WWW is control of subject interaction with the graphical user interface (GUI) of the experiment. The usage of the -BACK- and -FORWARD- buttons of current browsers is out of the experimenters control. Here, IT-based solutions have to be provided to prevent mis-usage of the GUI.

The motivation of subjects seems to be of considerable importance, because subjects might terminate participation at any time of the experiment. Situations, where subjects think they have to explain the interruption to the experimenter are unlikely to happen. The probability of drop-out in the Internet experiment seems higher than in the laboratory experiment (Reips 1997). Therefore, the experimenter should have control over the rate of general drop-out. Especially a selective drop-out should be traced, where subjects leave the experiment with different frequencies depending on experimental conditions (Reips 1997). Similar to this, subjects should take the recommended time to solve the problem.

Can Internet experiments satisfy requirements of experimental economics? Compared to laboratory experiments, the Internet experiment does not provide the experimenter with the same control. Additional noise is added: one source is the use of public networks, another is the lack of control of the subjects' environment, may that be at home or at work.

The Internet experiment seems to be less influenced by systematic errors, however, more random noise is added. Therefore, results produced by the new medium should be applicable in a more general way than laboratory results. This implies, that economic experiments could provide more parallelism to the field by using the benefits of the Internet.

3 IT-based solutions to problems identified

This section discusses techniques to solve the issues outlined in the methodological discussion in the previous section. In a first step techniques to increase subject control are discussed and in a second step the focus is on the control over the environment.

---

might enable participants to meet virtually even though they do not know themselves in person.

3.1 Techniques to increase subject control

The main objective with regards to subject control on the Internet is to uniquely identify subjects. Many online systems identify its participants via their e-mail addresses. Since it is not a problem to posses several different e-mail addresses, this is not a feasible approach for Internet experiments with the objective to avoid double participation by the same subject. Demanding very personal information, like social security number, may lead to self selection. A technique based on identification number or a social security number might not be effective, as there is no way to check whether a given number belongs to the participant.[8] Demanding very personal information from subjects may lead to self selection. For similar reasons do experiments conducted with a lottery-based reward ask for payment information only after the subject has been identified to win. In this case double-participation can only be determined at this (too) late stage.

To reduce the potential that one subject performs the same experiment under several different identities, payment information is used in addition to the e-mail address to uniquely identify each subject. While this reduces the number of potential double participants, it does not fully prevent such abuse. So far we use a traditional payment system, transfers to checking accounts, to identify subjects.

Currently we evaluate the integration of electronic payment systems in order to provide payments right after the experiment. For the above mentioned reasons, it is particularly important that the payment-system uniquely identifies the recipient of a transaction for Internet experiments.[9] In addition, the electronic payment mechanisms must operate timely and credit the necessary funds in a predictable manner. Still many electronic payment systems have been developed for different target applications in mind. Several systems provide extensive anonymity to the participants of a transaction and have properties similar to cash, e.g. DigiCash. These kind of systems do not provide the authentication needed for online economic experiments. Other systems, most notably credit card transactions via Internet, mainly specialized in customers paying their purchases. With regards to online experiments the option of peer-to-peer payments seem to be important, especially transfers from the experimenter to the subjects are necessary. Lately Paypal started to provide a successful peer-to-peer payment service that is widely used for online auction

---

[8] At least in Germany it is feasible to compute whether the personal identification number presented is a valid one since the algorithm is public.

[9] Subjects might not open up a new account exclusively for participation in the experiment since costs might outweigh the expected gains. One argument which threatens unique identification is the use of several payment accounts by a participant. In this case the experimenter can check for the same name.

payments. A good starting point on the evaluation of electronic payment systems can be found in MacKie-Mason & White (1997) and Schmidt & Müller (1999).

A Web based technique to augment subject identification is the use of "cookies". The cookie is an identifier that is stored on the users' client browser. Subjects are able to prevent their browser to accept cookies, they can remove them, or simply use a different web browser/computer. Therefore, a cookie is not a unique identifier of a subject. However, most Internet users are not familiar with these issues, so that this method can at least increase control.

The issue of whether or not a specific subject has fully independently performed an experiment cannot be avoided, since the experimenter has no control over the physical environment in which the subject is performing. The issue of one subject contacting another former subject to intentionally or accidentally pass on influential information can hardly be avoided for any experiment and only supported by organizational solutions like short availability of the experiment. Moreover, the chances of subjects actually knowing each other in a very diverse subject pool can be reduced by very selectively choosing among the user base of registered users, e.g., only notify one user per identifiable domain (company, department etc.). This is supported by presenting subjects individualized URLs, this means an identifier is added to the URL of the experiment site in order to make sure that only the invited subjects participate.

To cope with the issue of no-shows and subject's exact appearance, we are currently working on a reputation mechanism. In our system the subjects are rated with regards to their in time participation in former experiments. The experimenter can invite subjects on the basis of an index that distinguishes between positive and negative ratings.

3.2 Techniques to increase control over environment

There are several objectives with regards to control over the subjects' environment during the experiment. One demand of economic experiments is that only controllable help devices should be available. This includes the use of aids and tools, especially calculators. In case of Internet experiments only the use of provided aids can be controlled. Therefore, it seems important to integrate easy to use tools, like calculators, in the experiment to provide more control over the individual usage of helpers.

A further topic is to avoid and record drop-outs during the experiment. Most important is to distinguish between active and passive drop out. A passive drop out occurs due to a broken network connection or a crashed client computer. An active drop out occurs due to a user giving up. For either case it is easy to record the fact that for a specific user ID an experiment was not completed. Different options are available to handle this case. An experimentee may be given additional chances to complete the experiment on future logons, or maybe denied participation in the future.

It is more difficult to distinguish between both kinds of drop outs. To avoid active drop-out one needs to identify experimental stages with high drop-out rates and experiment with means to keep users' attention. Monetary rewards have to be calibrated in order to give subjects the right incentives not to drop out.

Participants should take the recommended time to solve the problem. We use several techniques to record time. Most important, the time is recorded at the beginning and the end of the experiment. In case economic decision variables are recorded in the database the time will be added as well. The interaction with the user-interface can be monitored by a standard log file of a web-server in case of a HTML based user interface. The only problem is to identify individual sessions when two different user use the same IP-address and browser.[10] Therefore, we propose a custom `access_log` which also includes the session-ID, the login of the user, and the experiment identifier.[11] The logfile is shown in Figure 1.

```
pD9E2695B.dip.t-dialin.net - - [21/Mar/2002:10:25:32 00100] "GET
/254551037374317/Welcome.mhtml HTTP/1.0" 200 5695

pD9E2695B.dip.t-dialin.net - - [21/Mar/2002:10:25:32 00100] "GET
/254551037374317/Welcome.mhtml HTTP/1.0" 200 5695 254551037374317
testuser AC
```

Fig. 1: Standard access_log top, custom access_log bottom

This custom logfile allows to compare actual and expected time to work over the whole and specific parts of the experiment of an individual subject identified by a username and a cookie based session-ID. And it provides a means of control of subject interaction with the GUI, i.e., which pages are requested.

A key design question for the implementation of the online experiment is the distribution of functionality to the server, to the client, or to both. It goes without saying that the implementation metaphor used must support this decision, i.e., pure HTML must be supported by server–side processing (e.g., JSP, PHP, or MetaHtml), Java-Script can realize client–side processing, and Java implements a hybrid approach. However, the more experiment functionality is "out sourced" from the server to the client, the higher the risk of potential

[10] This might be the case when two participants are connected to the experiment via the same proxy server.

[11] This function is actually implemented as a server side-include that writes the information including browser identification and referrer in a database. For a straight forward presentation the standard `access_log` file format is used in the figure.

fraudulent user interactions. The infrastructure provider has little control over the client side processing to prevent illegitimate use. On the other hand, data input integrity checks on the client side may substantially benefit the overall processing, since less browser–server interactions result, due to pre–validated input data (i.e., input errors are signaled to the user right away without a server connect).

Our experience in deploying economic experiments via the Internet has shown that some participants do — intentionally or unintentionally — try to "break" the system, e.g., through false input and aimless browser button use. This is a severe problem, since the experiment implementer has very little control over the way a participant uses the browser.

The solution we used for a previous experiment (Anderhub et al. 2001) is a client side approach by JAVA-applets and database connectivity classes (pre–JDBC) for the communication with the database. When using the client side approach, keeping state on the client is an easy task. In our case the database connectivity was implemented on the client, therefore some subjects could not participate because of policy restriction (e.g., proxy server and firewalls) of their local site connecting to the Internet. This problem can be eliminated by implementing a middle–tier, that handles the connection to the database, and communicates with the applet by standard http–protocol.

The server side approach to implement experiments uses (plain) HTML over the standard http–protocol; yet, it is more complex to keep track of the users' state. The technique we use in the current infrastructure is based on a finite state machine representation of the experimental stages. A stage constitutes a unit of interaction between the experimenter and the participant. This may include experiment instructions, decision forms, and questionnaires. Commonly an experimental stage corresponds to one page delivered to the user's browser (see Figure 2).

In this scheme one central file implements the finite state machine. It assembles the page transferred to the client browser dynamically according to current state and user action. This has two major advantages. One, the server maintains all state information and releases the right information (according to the experiment design) to the user. Two, a single URL is associated with this central control file, i.e., a user cannot jeopardize the operation of the experiment through intentional manipulation of the current URL (e.g., by guessing URLs).

Experimental stages are represented by states, and all possible user actions represent state transitions. A user action, for instance, is the pressing of a form–submit button, but also the forward or backward browsing through the experiment instructions. We implement all actions via form buttons, except for actions performed by the user with her browser (e.g., a page re–load). These latter actions constitute an issue, since at the server–side it may not al-

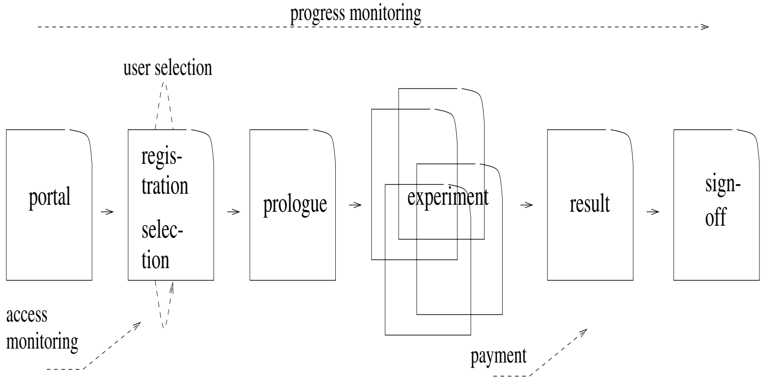

Fig. 2: Experimental stages

ways be possible to "catch" these actions (e.g., a back–button action might access the client browser page–cache only). We influence this by setting cache invalidate flags, so that a page that should not be re-loadable from the cache will be fetched from the server by the browser upon the occurrence of the re-load action. However, this requires the correct implementation of this protocol by the browser used.

To reduce the subject–browser interaction to the allowed actions, one can use the following approach. First, the experiment will open in a separate window by using JavaScript. Thus, JavaScript has to be enabled in order to participate. The experiment window does only contain the delivered HTML page, all buttons, address windows, title, and status bars are disabled. Second, in the experiment window the user is prevented from accessing the context-menu via the right mouse button by using the now surely enabled JavaScript again. At least, this is possible for the most used browsers Internet Explorer and Netscape Navigator. Thus, the only interaction a user can do with the experiment is either to use the provided buttons and links on the experiment page or to close the window and drop out. Although this method is not 100% safe, it increases control over the subject's browser interaction. Sample code is provided on our Web page.[12]

4 Infrastructure architecture and components

In this section we provide an overview of the design of the experiment infrastructure we have built to perform economic experiments on the World Wide

[12] `http://experiment.mpiew-jena.mpg.de/virtlab`

Web. Ultimately, we aim at offering an Internet service for use by the research community to perform and to participate in online economic experiments. This service offers a set of functions commonly needed by the experimenter, such as accounting, user authorization, and registration. Furthermore, our approach is to provide an environment that automates many tasks that have to be performed by the experimenter, such as participant selection and payment.

The infrastructure aims at providing maximal control over the experiment and its subjects to cope with the issues outlined in the previous sections. Figure 3 depicts the individual components of the environment described in more detail below.

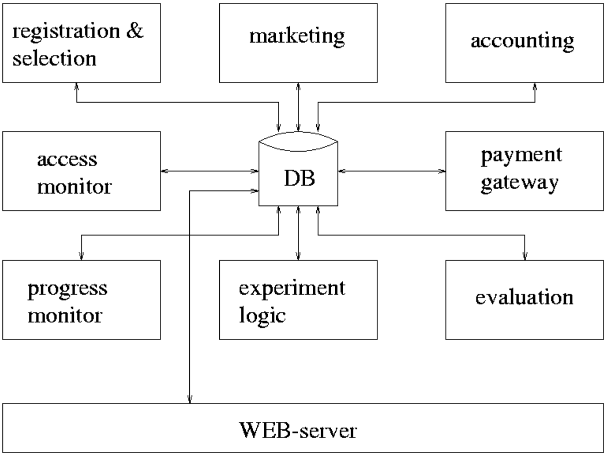

Fig. 3: Architecture of the experiment infrastructure.

The *registration* and *selection component* performs user registration and makes user selection decisions. Registration is a simple dialog asking the user to enter name, e-mail, banking account number, and other information. To ensure a certain degree of integrity of the entered data we interact with the user by sending her an access code by e-mail. The e-mail is generic not revealing the experiment's URL. A user must enter her identification and access code to actually sign–on for the experiment. With this procedure we want to ensure the presence of a valid communication channel with the subject.

User selection is done while registration is in progress. It is based on a set of rules granting or denying access for a user. Part of the rules derive from the particular needs of the experiment provider, who wants to address mostly students, or only females, for instance. Other rules are generic and directly address the experiment integrity, i.e., ensure that a novel subject is performing the experiment, for example. Note, that both kind of rules can only approximate the imposed constraints, and true integrity cannot be guaranteed in both cases. We think it is therefore best not to reveal the involved rules at this point.

The *access monitor* records access time, IP–address of client, system and browser type of client, and other information. The collected information is used by the selection sub–system, for instance, to derive access decision. The information is also used by the marketing component, to measure marketing efficiency and direct further advertising efforts.

The *progress monitor* records similar information during the entire experiment and ensures that parameters returned to the client are appropriately set (e.g., cache–reload attributes). We use a set of techniques to monitor the progression of the subject in the experiment. An analysis of this information may reveal that a subject went multiple times over the experimental instructions while being asked for decision values.[13] This information may be helpful in evaluating the individual results.

The *accounting monitor* component manages the financial assets for the experiment and communicates about payments with the *payment gateway*, the interface to the banking system connected. Our design aims at maximizing security and control over the financial assets available. At each moment in time the component knows exactly how much money is still available. If a critical limit is reached or an unusual high amount is being transferred an administrator is immediately notified and the experiment is halted. Default thresholds are defined and may be configured for each new experiment.

The *payment gateway* is the interface to the banking mechanism used. We aim at supporting several mechanisms: manual banking, online banking, and electronic payment systems, as they become available. The ultimate goal is to provide a fully automated payment system integration.

The key problem we are facing is the incapability to perform peer–to–peer payments via credit cards, the primary means of payment on the Internet to date. Unlike most electronic commerce transactions, the economic experiments we are targeting do require that the Internet service (the experimenter) pays the customer (the subject). Only recently banks seem to offer an online API (in Germany the standard HCBI is emerging) to effectuate customer transactions automatically. Until now, online banks offer their services through an HTML–based form, or Java–applet targeted at the *human user*. The lack of a standardized online banking API renders program–driven payment transactions very difficult.

The *experiment logic interface component* constitutes a set of interfaces that permits to plug in experiment implementations. The interface is open, and any implementation compliant to the interfaces may be *plugged in*. Tech-

[13] In case the experimental design foresees such actions.

niques for the management of methods in the Internet environment, which have been developed within the context of the MMM project (Jacobsen, Günther & Riessen 2000), are used to realize this component.

The *evaluation component* serves as direct interface to a statistics packages and to perform result evaluation "on–the–fly". This can be useful for standardized questionnaires provided by the experimental service. For specific experimental data the experimenter must identify how results are aggregated and evaluated, further processing of the results may then be carried out automatically. This evaluation is only a rough estimate and preliminary step, since outliers in the data are difficult to recognize automatically. We are currently working out details of this component, such as online vs. off-line processing and integration in the overall experiment infrastructure.

The *marketing* and recruitment component serves to advertise the experiment before and possibly during the experiment. It draws upon a large database of e-mail lists, individual e-mail addresses, newsgroups, and free Internet–ad space providers. The data inherent to this component is highly domain dependent, and will have to be carefully collected for alternate use. The component automates the sending of e-mail announcements to lists, the posting in news–groups, and the advertising of the experiment on free Internet–ad sites. The component also automates the return traffic processing as much as possible, e.g., management of bounced e-mails. The experimenter states in a graphical query what kind of subjects she wants to address: students or general public, specific sex, and/or geographical origin. At the end the experimenter states how many subjects should participate. Finally, the query draws the specified number of subjects randomly out of the eligible ones. In the future we aim at further developing the functionality of this component by also incorporating paid–ad providers and means to analyze feedback.

The *sign-off* component is a very simple component that manages mailing and interest lists. It prompts the user and, if she is interested, signs her up for different mailing lists concerning distribution of research reports about the experiment and further experimental economics research.

The implementation of the experiment infrastructure is based on Meta–HTML (Fox 1998), a server–side include programming environment that enables to establish and maintain session state, to manipulate databases out of HTML–documents, and to author dynamic HTML–pages, among others. The components are built around a database that maintains all experiment and participant data.[14] The Virtual Laboratory is online[15] and currently draws on a mailing list of more than 1,000 former participants.

[14]The infrastructure is build around open source software components: Meta-HTML Web server and scripting language, MySQL database.

[15]`http://experiment.mpiew-jena.mpg.de`

5 Conclusion

Internet experiments have become a popular tool for several research disciplines, such as experimental economics and experimental psychology. We have outlined several methodological constraints that govern Internet experiments as opposed to computer–based laboratory experiments. One of the major restrictions is the lack of control over the participant. We have motivated the design of an experiment infrastructure that aims at providing an improved degree of control and an automated management of many experiment tasks to the experimenter. The infrastructure we are developing constitutes a generic system with functional entities used in most e–commerce systems. These components comprise access monitoring, progress monitoring, marketing, user authorization and registration, and payment system integration. Furthermore these components may be used for online polling and market surveys, alike. In the future we aim at offering these infrastructure services to the research community to perform online experiments.

Acknowledgements

The financial support from the German Research Society (DFG grant nos. SFB 373/A3 and GRK 316) and the Max Planck Society is gratefully acknowledged. We would like to thank Werner Güth and two anonymous referees for valuable comments on a previous manuscript and comments from the audience at the Berlin Internet Economics Workshop 1999. We owe a great debt to coauthors and colleagues who helped us learn and improve Internet experiments, especially Vital Anderhub, Rudolf Müller, and Martin Strobel.

References

Anderhub, V., Müller, R. & Schmidt, C. (2001). Design and evaluation of an economic experiment via the Internet, *Journal of Economic Behavior and Organization* **46**: 227–247.

Baier, A., Bolle, F., Buschbaum, L. & Swiniarska, M. (1997). Please believe me! - Bargaining under incomplete information with communication, *Discussion Paper 98*, Fakultät für Wirtschaftswissenschaften Europa-Universität Viadrina, Frankfurt(Oder).

Birnbaum, M. H. (2000). Decision making in the lab and on the web, *in* M. H. Birnbaum (ed.), *Psychological Experiments on the Internet*, Academic Press, pp. 3–34.

Budimir, M. & Rieck, C. (1998). A portfolio selection experiment in the Internet, *Discussion paper*, Institut für Kapitalmarktforschung, Johann Wolfgang Goethe-Universität, Frankfurt am Main.

Camerer, C. (1997). Rules for Experimenting in Psychology and Economics, and Why They Differ, *in* W. Albers, W. Güth, P. Hammerstein, B. Moldovanu & E. van Damme (eds), *Essays in Honor of Reinhard Selten*, Springer, Berlin, pp. 313–327.

Drehmann, M., Oechssler, J. & Roider, A. (2002). Herding and Contrarian Behavior in Financial Markets - An Internet Experiment, *Bonn Econ Discussion Paper 25*, Bonn University.

Forsythe, R., Nelson, F., Neumann, G. & Wright, J. (1992). Anatomy of an Experimental Political Stock Market, *American Economic Review* **82**: 1142–1161.

Forsythe, R., Rietz, T. & Ross, T. (1999). Wishes, Expectations and Actions: A Survey on Price Formation in Election Stock Markets, *Journal of Economic Behavior and Organization* **39**: 83–110.

Fox, B. J. (1998). *The Meta-HTML Reference Manual*, Universal Access Inc. [web page] URL http://www.metahtml.org.

Friedman, D. & Sunder, S. (1994). *Experimental Methods - A primer for Economists*, Camebridge University Press.

Holt, C. (2002). Webgames and Strategic Behavior: Recipes for Interactive Learning, *mimeo*.

Jacobsen, H.-A., Günther, O. & Riessen, G. (2000). Component Leasing on the World Wide Web, *Netnomics* **2**(2): 191–219.

Kirchkamp, O. (2000). WWW Experiments for Economists, a Technical Introduction, *Discussion Paper 00-09*, Sonderforschungsbereich 504, Universität Mannheim. [web page] URL http://www.sfb504.uni-mannheim.de/~oliver/webexp.html.

Krantz, J. H., Ballard, J. & Scher, J. (1997). Comparing the results of laboratory and World-Wide Web samples on the determinants of female attractiveness, *Behavior Research Methods, Instruments, & Computers* **29(2)**: 264–269.

Lucking-Reiley, D. (1999). Using field experiments to test equivalence between auction formats: Magic on the internet, *American Economic Review* **89**(5): 1063–1080.

MacKie-Mason, J. & White, K. (1997). Evaluating and selecting digital payment mechanisms, *in* G. L. Rosston & D. Waterman (eds), *Interconnection and the Internet. Selected Papers from the 1996 Telecommunications Policy Research Conference*, Mahwah, NJ: Lawrence Erlbaum [pdf-document] URL http://www-personal.umich.edu/~jmm/papers/digidoll.pdf.

Reips, U.-D. (1997). Das psychologische Experimentieren im Internet, *in* B. Batinic (ed.), *Internet für Psychologen*, Hogrefe-Verlag, pp. 245–265.

Schmidt, C. & Müller, R. (1999). A framework for micropayment evaluation, *Netnomics* **1**(2): 187–200.

Shavit, T., Sonsino, D. & Benzion, U. (2001). A comparative study of lotteries-evaluation in class and on the web, *Journal of Economic Psychology* **22**: 483–491.

Smith, V. L. (1976). Experimental economics: Induced value theory, *American Economic Review* **66**(2): 274–279.

Wurman, P. R., Wellmann, M. P. & Walsh, W. E. (1998). The Michigan Internet AuctionBot: A configurable auction server for human and software agents, *Proceedings of the Second International Conference on Autonomous Agents*, Minneapolis, MN, USA.

TRAX
A Program Package for *in vivo* Visualization of Human Brain Fibers

Benno Pütz, Kemal Ates, Dorothee P. Auer
Max-Planck-Institute of Psychiatry, München

Abstract

Recent developments in magnetic resonance imaging (MRI) have made it possible to detect nerve fiber bundles in the white matter of the human brain. While this information is highly valuable in, e. g., pre-surgical planning, new visualization facilities had to be developed to adequately tap this source of information.

We have, thus, developed an application that aims at maximizing the amount of information for the user by combining anatomical and functional MR images with nerve fiber bundles, derived from diffusion tensor imaging, in an interactive 3D visualization that is intuitive to understand and use. This new tool has obvious applications in the field of pre-surgical planning, but also in a more general sense for studies of brain connectivity and function. Due to its interactive interface it should also be an ideal tool in medical teaching.

1 Introduction

Until recently, the structure of the living brain was accessible only through more or less non-invasive imaging techniques which allowed essentially the differentiation of grey and white matter. In many cases magnetic resonance imaging (MRI) is the technique of choice as it combines the advantage of no known negative effects with a large arsenal of acquisition techniques that allow the imaging of various aspects other than structure alone, e. g., perfusion or function-related brain activation.

Non-invasive functional mapping techniques such as functional magnetic resonance imaging (fMRI) as well as electro- or magneto-encephalogramms (EEG/MEG), provide further information about the brain's cortical organization, i. e., they help identifying the cortical areas associated with the functions investigated (e. g., motor control, speech). This kind of information has greatly helped further the knowledge about the function of the brain. For neurosurgeons this knowledge is essential for planning in order to spare vital areas and, thus, their functions, during surgery. All of this is well-established and used routinely in medical centers world-wide.

This kind of map, however, only describes the outermost 4–6 mm thick layer of the brain, the cortex (also called *grey matter*). In contrast, until recently, the remaining parenchyma, the underlying white matter, could not be differentiated any further. White matter consists of axonal connections either between cortical areas, subcortical nuclei, or to the spinal cord. It is, therefore, also important to know the location of the fiber tracts interconnecting the areas identified by, e. g., fMRI, since they are equally essential for the preservation of the corresponding function.

At this point the relatively novel technique of *diffusion tensor imaging* (DTI), introduced by Basser and Pierpaoli (1996), in combination with *fiber tracing*, first demonstrated by Mori *et al.* (1999) and Conturo *et al.* (1999), comes into play. It is based on the fact that the diffusive mobility of water molecules is higher along an axonal fiber than perpendicular to it. Using appropriate acquisition schemes it is possible to assess that mobility and derive a simple model described by the so-called diffusion tensor. From this model it is possible to infer the local direction of axonal fibers and, thus, to trace trace fiber bundles.

Over the last year a system has evolved that originally targeted the needs for 3D visualization of DTI-derived fiber information in an attempt to better understand the brain's connectivity.

The result is a very general visualization tool for all kinds of medical image data that has many applications in a much broader range than initially intended, from basic neurosciences, to teaching tool for neuroanatomy, and planning aid for neurosurgery.

2 Background

Conventional brain MRI data provide two-dimensional slices through the head, which may be stacked to obtain three-dimensional coverage. The images represent some scalar parameter (e. g., relaxation time) acquired on the regular grid of the image voxels and are usually displayed as gray-scale image.

DTI provides a more complex parameter, th so-called diffusion tensor. The underlying idea is that the diffusive mobility of water molecules (which are the ultimate signal source in MRI) may depend on their surroundings. In white matter, where fiber bundles are formed by many axons running parallel, the mobility along the axons is higher than perpendicular to them[1], we observe anisotropic diffusion. With an appropriate acquisition sequence of diffusion-weighted measurements this difference can be detected and approximately described by a diffusion tensor field (Basser and Pierpaoli (1996)). Mathematically the diffusion tensor is a symmetric 3×3 matrix, and, as such, is not easily understood or visualized. There is, however a geometric interpretation of its eigenvectors and -values as the principal axes and the corresponding diffusivity of the related diffusion ellipsoid.

In a nerve fiber bundle, the eigenvector corresponding to the largest eigenvalue can, thus, be interpreted as the local fiber orientation. It follows that the orientation field of the white matter fibers can be obtained by a diagonalization of the diffusion tensor field. This step reduces the complexity of the observed parameter to a vector field, which can either either be visualized directly (e. g., RGB-coded) or taken as the basis for fiber reconstruction algorithms that try to fit fibers into the direction field. It is the visualization of this kind of data, no longer restricted to the grid of the image voxels but continuous in the acquisition space, that motivated the development of TRAX, as there was no software available for this task, not even commercially.

It is possible to extract scalar quantities from the diffusion tensor, further reducing the data complexity. Some meaningful quantities of this type are mean diffusivity and anisotropy measures.

3 Design

From that starting point we tried to implement a package that would handle all not only the “conventional” display of MR data (anatomic and functional) in either two or three dimensions but also arbitrary objects placed in the brain.

[1] In principle this holds in grey matter, too, however the extent of parallel fiber pieces there is much smaller than the voxel dimension. The effect is averaged away by many more or less random fiber orientations for any pixel containing grey matter.

3.1 Design Goals

3.1.1 Ease of Use

In order to be used by a wide range of targeted users, each component had to designed to used intuitively.

3.1.2 Functionality

The two main function blocks are an interface to various intermediate modules that generate data such as fiber definitions or segmentations and the visualization

Visualization

– fibers
– anatomical (high-resolution) MR images
– additional information (e. g., from fMRI analysis)

3.1.3 Modularity

As some components have already been available and this has been a collaborative endeavor modularity was one major design goal. It should be easy to simply plug-in functional modules developed elsewhere.

This proved also advantageous as we work in a rapidly evolving field in the application of MR imaging, where improvements to various stages of the data processing stream are still proposed frequently. An open architecture allowed for easy inclusion of new modules to adapt the systems capabilities.

This also applies to the acceptance of various data formats which is still being extended. As many groups who acquire DTI image data have their own code for calculating the diffusion tensor without any agreed standard format, the program understands some common formats and is easily extended to read different ones.

3.2 Language

For the implementation of TRAX we choose IDL (Interactive Data Language, http://www.rsinc.comResearch Systems, Inc.) as a high level programming language to keep development times short. Much of the previous data processing code, especially input/output and format conversion routines, but also the DTI calculation, was already written in IDL and provided a basis to start with. Its object-oriented graphics system provides excellent support for three-dimensional graphics with reasonable performance on current LINUX workstations. This was one of the main reasons for IDL, as we do not have the personel to reimplement all those capabilities ourselves. The add-on IDL-on-the-Net (ION) would allow access to the system through WEB browsers, an option that might be investigated more thoroughly in the future. The downside of our choice is that IDL is only available commercially.

Some of the plug-in modules were written in collaboration with other groups using either C++ or Fortran90. To allow module development without requiring IDL running at the development site and to be independent of language-specific differences in data passing, the modules are written as command line tools, using temporary files as interface to TRAX.

Some recent utility routines have been implemented in JAVA, most notably the generation routines for animated movies. In the future this may be incorporated in a WEB-based interface, allowing remote sites to gain access to precompiled results.

4 Description

The two main function blocks are region definition/fiber tracing and the interactive display.

4.1 Data preparation

Fiber tracing is performed by placing seed points in the orientations field derived from the diffusion tensor and then following the local direction. Several different approaches to this problem have been proposed in the literature (Conturo *et al.* (1999); Mori *et al.* (2000); Gössl *et al.* (2002)). Keeping the system modular, we decided to keep that part in the form of separate subprograms that would be called from the interface after defining appropriate starting regions.[2]

Start regions are outlined in a region definition module where MR data serve as template.

As the choice selection of what might be appropriate starting regions is not always possible from a structural image alone, various image modalities can be chosen from, including functional activation and derived maps from DTI acquisition such as anisotropy and weighted diffusion direction. It is also possible to use one modality as background and overlay another semi-transparently, as indicated in the inlays of Fig. 1.

The outlines are saved as binary masks in the same format the MR images are stored. They can, therefore, be included in the visualization routines,

[2] A solution for the start-end problem is currently being worked on and already being addressed in the interface.

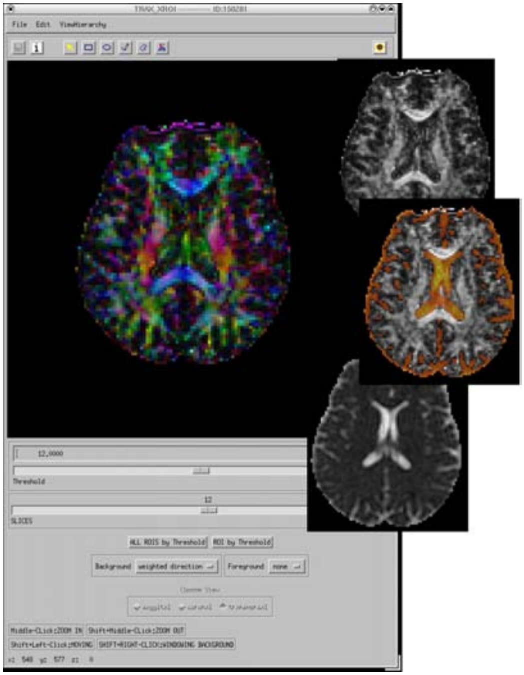

Fig. 1: ROI definition interface with various overlay modi demonstrated in inlays.

displayed as either surface or volume (see below). This proved useful for manual segmentations of various brain regions that had to be shown in their relative position to fibers.

4.2 Fiber tracing

For the determination of fiber tracts we currently employ the algorithm published by Gössl *et al.* (2002).

4.3 Visualization

The main focus of our project has been the usability of the technique in a clinical environment. To this end the data had to be visualized in a way that made interpretation intuitive, even to an untrained user. An example is shown in Fig. 2.

The basic display consists of a fiber reconstruction, i. e., all the fiber lines that resulted from the starting region(s) of one tracing run, within the corresponding brain surface. However, arbitrarily complex configurations can

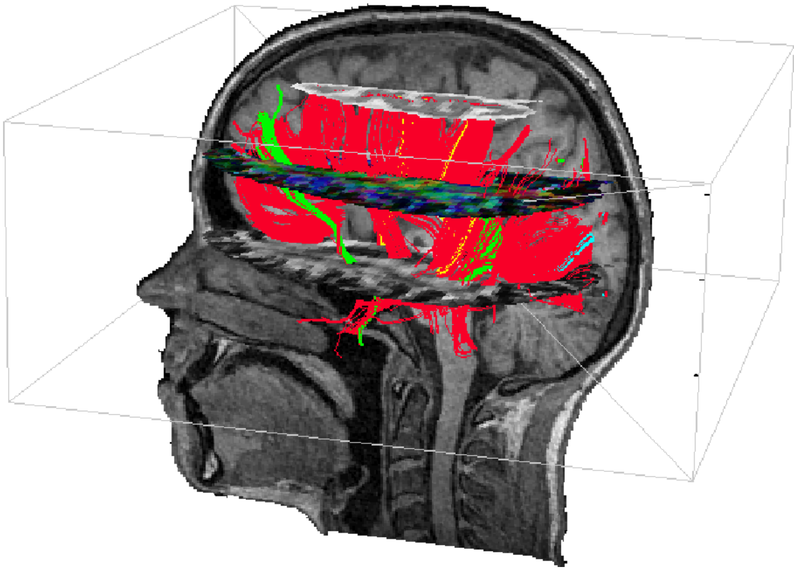

Fig. 2: Display of trans-collossal fibers showing some of the program's capabilities: highlighting selected fibers, slice displays of various modi (T_1, diffusivity, anisotropy, fiber orientation).

be constructed by adding other objects. These might be other fibers or segmented subvolumes from the DTI data set as, e. g., the tumor show in Fig. 4. It is also possible to fuse data derived from other acquisitions, e. g., high-resolution T_1 or T_2 weighted images. From such data sets one might extract the skin surface or simply slices that can be placed in the view.

In addition, 2D-slice views similar to those routinely used by clinicians can be generated with overlays for the fibers. Examples can be seen in Fig. 3.

While the interactive manipulation of the view gives a good impression of the relative positions of objects in space, stereo views offer a way to provide as much of that information as is conveyable in print. An example is shown in Fig. 4.

4.4 Modality fusion

Larger functional networks, e. g., the language system, are constituted from several spatially separated areas which, working together, achieve the network's combined performance. Understanding the way these areas are interconnectedmay shed some light on how the brain realizes some of its more complex functions.

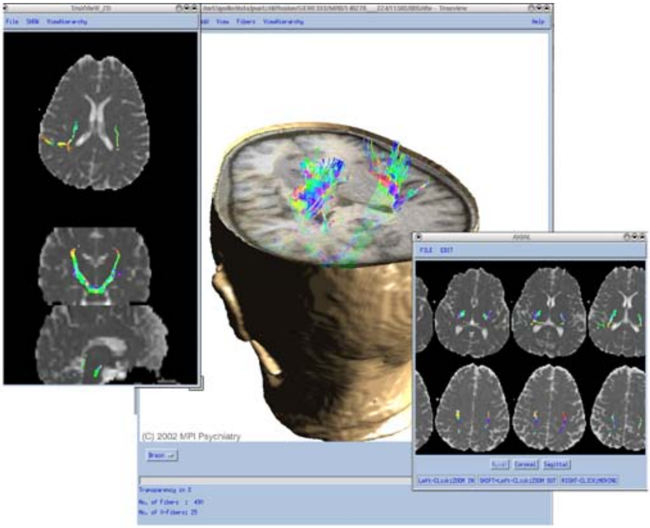

Fig. 3: Complex object composition with anisotropy-colored fibers in 3D and the two 2D viewing modes.

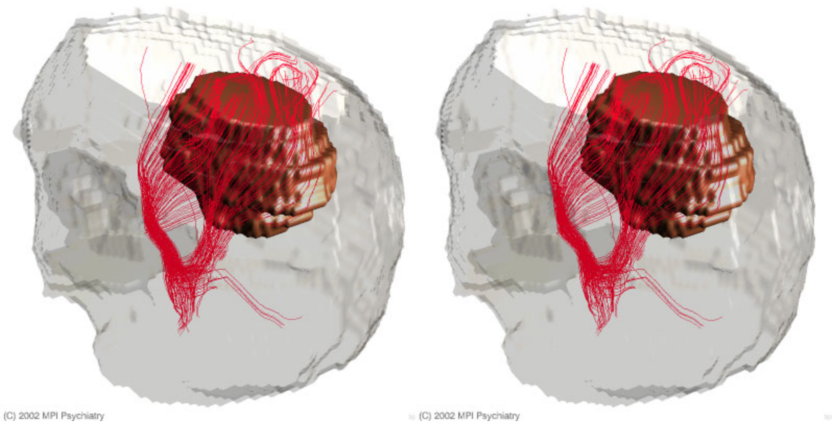

Fig. 4: Stereo view of fibers displaced by a large temporal tumor.

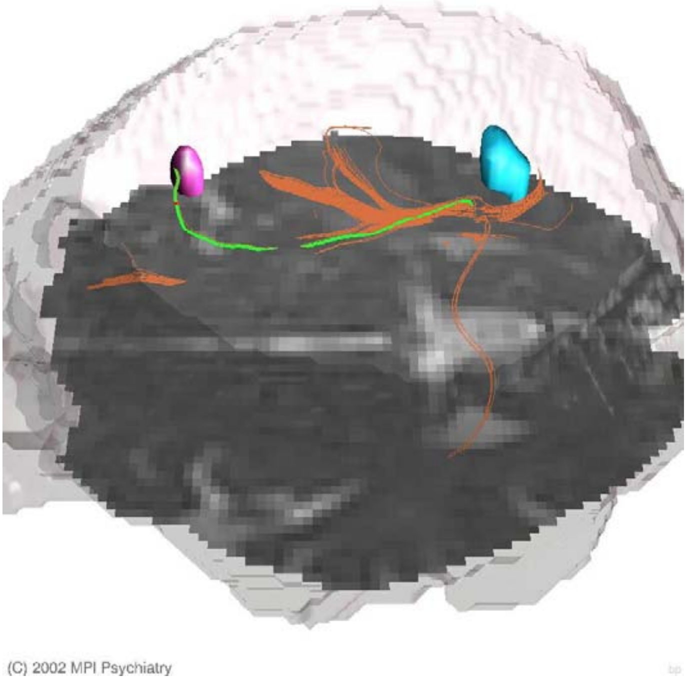

Fig. 5: Combination of fMRI and DTI to investigate the connectivity in the language areas. The two blobs indicate Broca's and Wernicke's area, respectively, the highlighted fiber is probably part of the arcuate fascicle.

In order to understand the brain's connectivity, one starts with the delineation of functional areas in the cortex. This task is usually performed with fMRI, providing in so-called activation maps, that indicate areas involved with certain brain activities. In functional networks one expects fiber connections between related areas, either directly or via some intermediate nodes.

It was, thus, an essential requirement to be able to fuse functional activation maps with DTI-derived fibers to reveal their respective spatial arrangement.

5 Applications

TRAX and its predecessors have been in use for in-house patient monitoring for over a year now. Especially tumor patients undergo a DTI examination with subsequent fiber tracing. An example is shown on Fig. 4 where a patient with a large tumor is depicted in a stereo view.

Nonetheless,

6 Outlook

A system has evolved that proved capable of delivering the envisioned functionality in a research environment. Due to its open structure and modular design, it is easily extended to new data formats as well as new algorithms in either the data processing stream or the visualization.

First results have been published in conference proceedings (e.g., Auer *et al.* (2001); Pütz and Auer (2002)) and journals (Gössl *et al.* (2002)). First external collaborations have been initiatedand further requests for have been received after presentations at workshops and conferences.

We continue to improve and extend the system:

- include more data formats / remove site-dependence
- eventually allow the inclusion of different modalities (PET, EEG/MEG)
- improve the remote access capabilities
  - Simple result presentation (animations)
  - Relatively complete access to calculation facilitites (ION)
- more seamless integration of external modules

Acknowledgements

TRAX has been made possible by the continued support Ludwig M. Auer who initiated the engagement in pre-surgical DTI. We would also like to thank Daniel Boos for his assistance in writing early versions of the code, and Christoff Gössl, Uli Freudensprung, Susanne Heim, Ludwig Fahrmeir, and Klaus Hahn for contributing some of the external plug-ins (fiber tracing and non-linear smoothing) necessary for generating good fiber estimates.

References

Dorothee P. Auer, Benno Pütz, and Ludwig M. Auer. Visualization of major fiber tracts and cortical functions in patients with cerebral gliomas. In *Proc. of the 2nd Int. Mt. Bandai Symposium for Neuroscience*, pages 115–119. Ishii Densan Printing Co., Ltd., Koriyama, Japan, 2001.

Peter J. Basser and Carlo Pierpaoli. Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI. *J. Magn. Reson.*, 111(3):209–219, 1996.

Thomas E. Conturo, Nicolas F. Lori, Thomas S. Cull, Erbil Akbudak, Abraham Z. Snyder, Joshua S. Shimony, Robert C. McKinstry, Harold Burton, and Marcus E. Raichle. Tracking neuronal fiber pathways in the living human brain. *Proceedings of the National Academy of Sciences of the United States of America*, 96(18):10422–10427, August 1999.

Christoff Gössl, Ludwig Fahrmeir, Benno Pütz, Ludwig M. Auer, and Dorothee P. Auer. Fiber tracking from DTI using linear state space models: Detectability of the pyramidal tract. *NeuroImage*, 16(2):378–388, June 2002. DOI: 10.1006/nimg.2002.1055.

Susumu Mori, Barbara J. Crain, V. P. Chacko, and Peter C. M. van Zijl. Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Annals of Neurology*, 45(2):265–269, February 1999.

Susumu Mori, Walter E. Kaufmann, Godfrey D. Pearlson, Barbara J. Crain, Bram Stieltjes, Meiyappan Solaiyappan, and Peter C. M. van Zijl. In vivo visualization of human neural pathways by magnetic resonance imaging. *Annals of Neurology*, 47(3):412–414, March 2000.

Benno Pütz and Dorothee P. Auer. Towards validation of DTI-based fiber tracking: Plant phantoms. In *NeuroImage Human Brain Mapping 2002 Meeting*, page 487, 2002.

Anschriften der Autoren

Kemal Ates
Max-Planck-Institut für Psychiatrie
Kraepelinstr. 2 und 10, 80804 München
E-Mail: kemal@mpipsykl.mpg.de

Dorothee P. Auer
Max-Planck-Institut für Psychiatrie
Kraepelinstr. 2 und 10, 80804 München
E-Mail: auer@mpipsykl.mpg.de

Daan Broeder
Max-Planck-Institut für Psycholinguistik
Wundtlaan 1, 6535 XD Nijmegen, Niederlande
E-Mail: daan.broeder@mpi.nl

Hennie Brugman
Max-Planck-Institut für Psycholinguistik
Wundtlaan 1, 6535 XD Nijmegen, Niederlande
E-Mail: hennie.brugman@mpi.nl

Reiner Dirksmeyer
Max-Planck-Institut für Psycholinguistik
Wundtlaan 1, 6535 XD Nijmegen, Niederlande
E-Mail: reiner.dirksmeyer@mpi.nl

Thomas Fischbacher
Max-Planck-Institut für Gravitationsphysik,
Albert-Einstein-Institut
Mühlenberg 1, 14476 Potsdam
E-Mail: tf@aei.mpg.de

Ben Greiner
Max-Planck-Institut zur Erforschung von Wirtschaftssystemen
Kahlaische Str. 10, 07745 Jena
E-Mail: greiner@mpiew-jena.mpg.de

Hans-Arno Jacobsen
Department of Electrical and Computer Engineering and
Department of Computer Science
University of Toronto
E-Mail: jacobsen@eecg.toronto.edu

Dr. Frank Jenko
Max-Planck-Institut für Plasmaphysik
Boltzmannstr. 2, 85748 Garching
E-Mail: fsj@ipp.mpg.de

Benno Pütz
Max-Planck-Institut für Psychiatrie
Kraepelinstr. 2 und 10, 80804 München
E-Mail: puetz@mpipsykl.mpg.de

Carsten Schmidt
Max-Planck-Institut zur Erforschung von Wirtschaftssystemen
Kahlaische Str. 10, 07745 Jena
E-Mail: cschmidt@mpiew-jena.mpg.de

In der Reihe GWDG-Berichte sind zuletzt erschienen:

Nr. 40 *Plesser, Theo und Peter Wittenburg* (Hrsg.):
Forschung und wissenschaftliches Rechnen – Beiträge zum Heinz-Billing-Preis 1994
1995

Nr. 41 *Brinkmeier, Fritz* (Hrsg.):
Rechner, Netze, Spezialisten. Vom Maschinenzentrum zum Kompetenzzentrum – Vorträge des Kolloquiums zum 25jährigen Bestehen der GWDG
1996

Nr. 42 *Plesser, Theo und Peter Wittenburg* (Hrsg.):
Forschung und wissenschaftliches Rechnen – Beiträge zum Heinz-Billing-Preis 1995
1996

Nr. 43 *Wall, Dieter* (Hrsg.):
Kostenrechnung im wissenschaftlichen Rechenzentrum – Das Göttinger Modell
1996

Nr. 44 *Plesser, Theo und Peter Wittenburg* (Hrsg.):
Forschung und wissenschaftliches Rechnen – Beiträge zum Heinz-Billing-Preis 1996
1997

Nr. 45 *Koke, Hartmut und Engelbert Ziegler* (Hrsg.):
13. DV-Treffen der Max-Planck-Institute – 21.-22. November 1996 in Göttingen
1997

Nr. 46 **Jahresberichte 1994 bis 1996**
1997

Nr. 47 *Heuer, Konrad, Eberhard Mönkeberg und Ulrich Schwardmann:*
Server-Betrieb mit Standard-PC-Hardware unter freien UNIX-Betriebssystemen
1998

Nr. 48 *Haan, Oswald* (Hrsg.):
Göttinger Informatik Kolloquium – Vorträge aus den Jahren 1996/97
1998

Nr. 49 *Koke, Hartmut und Engelbert Ziegler* (Hrsg.):
IT-Infrastruktur im wissenschaftlichen Umfeld – 14. DV-Treffen der Max-Planck-Institute, 20.-21. November 1997 in Göttingen
1998

Nr. 50 *Gerling, Rainer W.* (Hrsg.):
Datenschutz und neue Medien – Datenschutzschulung am 25./26. Mai 1998
1998

Nr. 51 *Plesser, Theo und Peter Wittenburg* (Hrsg.):
Forschung und wissenschaftliches Rechnen – Beiträge zum Heinz-Billing-Preis 1997
1998

Nr. 52 *Heinzel, Stefan und Theo Plesser* (Hrsg.):
Forschung und wissenschaftliches Rechnen – Beiträge zum Heinz-Billing-Preis 1998
1999

Nr. 53 *Kaspar, Friedbert und Hans-Ulrich Zimmermann* (Hrsg.):
Internet- und Intranet-Technologien in der wissenschaftlichen Datenverarbeitung – 15. DV-Treffen der Max-Planck-Institute, 18. - 20. November 1998 in Göttingen
1999

Nr. 54 *Plesser, Theo und Helmut Hayd* (Hrsg.):
Forschung und wissenschaftliches Rechnen – Beiträge zum Heinz-Billing-Preis 1999
2000

Nr. 55 *Kaspar, Friedbert und Hans-Ulrich Zimmermann* (Hrsg.):
Neue Technologien zur Nutzung von Netzdiensten – 16. DV-Treffen der Max-Planck-Institute, 17. - 19. November 1999 in Göttingen
2000

Nr. 56 *Plesser, Theo und Helmut Hayd* (Hrsg.):
Forschung und wissenschaftliches Rechnen – Beiträge zum Heinz-Billing-Preis 2000
2001

Nr. 57 *Hayd, Helmut und Rainer Kleinrensing (Hrsg.)*
17. und 18. DV-Treffen der Max-Planck-Institute,
22. - 24. November 2000, 21. – 23. November 2001 in Göttingen
2002

Nr. 58 Macho*, Volker und Theo Plesser* (Hrsg.):
Forschung und wissenschaftliches Rechnen – Beiträge zum Heinz-Billing-Preis 2001
2003

Nr. 59 *Suchodoletz, Dirk von*:
Effizienter Betrieb großer Rechnerpools – Implementierung am Beispiel des Studierendennetzes an der Universität Göttingen
2003

Nr. 60 *Haan, Oswald (Hrsg.)*:
Erfahrungen mit den IBM-Parallelrechnersystemen RS/6000 SP und pSeries690
2003

Nr. 61 *Rieger, Sebastian*:
Streaming-Media und Multicasting in drahtlosen Netzwerken – Untersuchung von Realisierungs- und Anwendungsmöglichkeiten
2003

Nr. 62 *Kremer, Kurt und Volker* Macho (Hrsg.):
Forschung und wissenschaftliches Rechnen – Beiträge zum Heinz-Billing-Preis 2002
2003

Nähere Informationen finden Sie im Internet unter
`http://www.gwdg.de/forschung/publikationen/gwdg-berichte/index.html`