40th

Cray User Group Conference

STUTTGART 1998

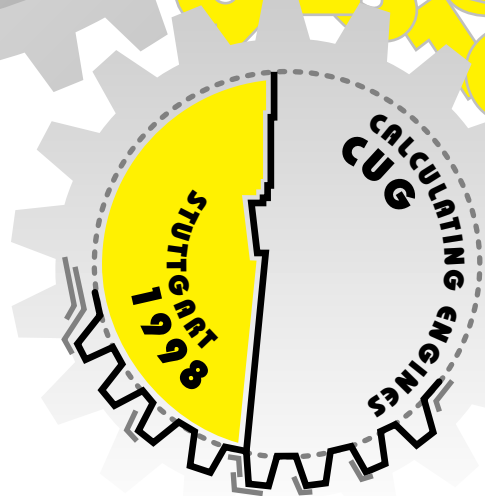CALCULATING CUG ENGINES

Final Program

June 15–19, 1998

Stuttgart, Germany

Welcome

Dear CUG members and friends,

It is with great pleasure that we welcome you to the 40th Cray User Group Conference. CUG 98 will provide you with the opportunity to hear the latest information about what's "hot" in High Performance Computing and Communication on SGI and Cray systems. debis Systemhaus, the Computer Center of the University of Stuttgart (RUS) and SGI Germany are pleased to jointly host this conference. The conference site, Kultur und Kongresszentrum Liederhalle (KKL), is nicely located in downtown Stuttgart within walking distance of many attractive places.

Our theme, "Calculating Engines," spans the days of Wilhelm Schickard who in 1623 designed in nearby Tübingen the first mechanical calculating engine (on display at the conference) through the famous Cray-2 at RUS, first outside the US, to today's High-Performance Computing Center, jointly operated by RUS and Daimler-Benz InterServices (debis) AG, the services company of the Daimler-Benz Group.

During CUG 98 you will meet many people with interests similar to yours. Visit beautiful Baden-Württemberg in the South of Germany in the early summertime. Stroll through the city of Stuttgart, the capital of Baden-Württemberg, surrounded by more hills than the seven of Rome, where you will find interesting and beautiful museums, theaters, and estates to visit—and have the chance to do some shopping in a modern city with exciting shops. Take a tour to the famous old university towns of Heidelberg and Tübingen not far away. See Lake Constance and the Black Forest. Drive the Baroque Route of the Upper Swabia. Visit the birthplaces of Schiller, Hegel, Mörike, Daimler, or

Bosch to name only a few famous persons from our region. Throughout your stay and travels, you will certainly enjoy our regional cuisine, in addition to that inspired by our French neighbors. You will find our famous wines to be a pleasure. Come see and feel the old culture of Europe.

To make the most of your stay with us, we have prepared this booklet which provides information about:
- Conference program,
- Conference registration,
- Conference hotels, and
- Social events.

Grüß Gott! in Stuttgart, the site of "Calculating Engines"!

Prof. Dr.-Ing. Roland Rühle
Director
Computer Center University of Stuttgart (RUS)

Karl-Heinz Streibich
Member of the Board
debis Systemhaus GmbH

Dipl.-Ing. Walter Wehinger
RUS
Local Arrangements Chair

R U S   debis Systemhaus

Message from Sam Milosevich, CUG Vice-President and from the CUG Program Committee

To the members of CUG:

Welcome to the 40th Cray Users Group meeting! This is your best chance in 1998 to exchange professional information and enjoy personal interactions with your fellow SGI/Cray Research supercomputing users.

While visiting the fabled capital of Baden-Württemberg, nestled in the south of Germany, you will discover how CUG—your SGI/Cray user forum for high performance communications and computing—can give you up-to-the-minute information and insight for that competitive edge you need.
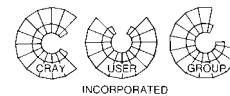
Please take a look at the technical program in the following pages. Here, you will find a wealth of great opportunities on the CUG theme of "Calculating Engines"—from the SGI/Cray kind during the day to the Mercedes-Benz kind for our traditional conference "CUG Night Out" event.

The CUG Program Committee has assembled for you a diverse array of detail-packed presentations in Tutorials, Keynote and General Sessions, Special Interest Technical presentations, and spontaneous Birds of a Feather (BOF) discussions.

- *Tutorial Sessions on Monday morning—available to you at no additional charge—are a great opportunity for you to update your technical skills with the help of selected technical experts from SGI/ Cray Research and CUG sites.*

- *The Welcome and Keynote Session, Monday afternoon, will be the Initial Program Load of "Calculating Engines" for CUG in Stuttgart, hosted by the Computer Center of the University of Stuttgart (RUS) and by Daimler-Benz InterServices (debis) AG.*

- *Parallel Technical Sessions each day will give you the opportunity to focus on the specific knowledge domains of the Special Interest Groups (SIGs). Presentations in these sessions have been reviewed and selected by the SIG chairpersons.*

- *Birds of a Feather (BOF) Sessions, the most dynamic aspect of a CUG conference, are scheduled as needed, and notices of BOF meetings will be posted near the Message Board. You are welcome to organize a BOF session at the conference.*

- *Formal receptions and informal luncheons are among the countless occasions you will have to exchange information with your colleagues from other CUG sites and to collaborate with representatives from SGI/Cray Research.*

You can see how the Stuttgart CUG offers something for everyone during an exciting, educational, and entertaining week of "Calculating Engines."

Sam Milosevich
CUG Vice-President and Program Chair
Eli Lilly and Company

Program Notes

Contents

Special Interest Forums

The Special Interest Groups (SIGs) hold meetings that are open to all interested CUG attendees. Depending on the SIG, these meetings take place on Monday or Thursday. They provide a forum to discuss the future direction of the SIG, important issues, and talks of special interest. You are encouraged to attend any of these open meetings.

Program Committee Meeting

The Program Committee needs your ideas! Come to the Open Program Committee meeting on Thursday to hear about the next CUG conference and to share your suggestions of how to make each CUG the best one to date. All conference attendees are invited and encouraged to attend this open meeting. The CUG Conference Program is the net result of the combined efforts of all interested CUG sites. Come to the Program Committee meeting and make your unique contribution to the next CUG conference!

Video Theater

Supercomputers generate data that, via scientific visualization, can be transformed into pictures that provide a way for scientists to understand the information more clearly and quickly. Although the primary purpose of visualization is to facilitate the discovery process, scientists are increasingly coming to rely on it also to present their research conclusions. This Friday morning session showcases the latest in scientific/engineering visualization.

Conference Program Updates

We anticipate that there may be continuing changes to this Program schedule. You should plan to check the Bulletin Boards on site at CUG each day for schedule changes. Thank you for your patience and cooperation.

Hohenzollern (4.4.22)

Tutorial I

8:30   UNICOS/IRIX Differences
in DMF
Paul Ernst, SGI/Cray

Rhein (4.3.11)

Tutorial II

8:30   Scheduling and Configuration for Origin 2000
Jim Harrell, SGI/Cray

Monrepos (4.4.25)

Tutorial III

8:30   Scheduling and Configuration Tuning for the T3E
Jim Grindle, SGI/Cray

Schiller-Saal

Tutorial IV

8:30   Visualization:
from Theory to Practice
Cal Kirchoff, SGI/Cray

10:30–11:00   Break

Tutorial V

11:00   Contemporary Issues in
Network Security
Jay McCauley, SGI/Cray

Tutorial VI

11:00   Performance Optimization
for the Origin 2000
Jeff Brooks, SGI/Cray

Tutorial VII

11:00   Mathematical Libraries
Bill Harrod, SGI/Cray

12:30–2:00   Newcomers Luncheon

Schiller-Saal

| **General Session**

Chair: Walter Wehinger, RUS

2:00 Welcome —University of Stuttgart
Prof. Dr. rer. nat. E. Messerschmid,
Prorektor Research and Technology

2:15 Parallel Numerical Simulations of
Environmental Phenomena
Prof. Dr. rer. nat. Gabriel Wittum, Director, Institute for
Computer Applications, University of Stuttgart

2:55 Challenges in the Digital World of Automotive R&D
Prof. Dr. Bharat Balasubramanian

3:25 CUG Report
Gary Jensen, CUG President, UIUCNCSA
Sam Milosevich, CUG Vice-President, ELILILLY

3:45–4:15   Break

Schiller-Saal

4:15–5:00   Future of CUG

Hohenzollern (4.4.22)	**Rhein (4.3.11)**	**Monrepos (4.4.25)**	**Solitude (4.4.20)**
2A Open Meetings	**2B** Open Meetings	**2C** Open Meetings	**2D** Open Meetings

5:00–6:00   Please plan to attend the **Open Meetings** to discuss topics of interest to each Special Interest Group.

Information on location and topics will be posted on the bulletin board on Monday, June 15.

SGI/Cray Research Reception (Alte Reithalle) 7:00 p.m.–10:00 p.m.

Hohenzollern (4.4.22)

3A The Engine's Core: Operating Systems

Chair: Ingeborg Weidl, MPG

9:00 PVP UNICOS Status and Update
Patti Langer, SGI/Cray

9:30 T3E and UNICOS/mk Status and Update, *Jim Grindle, SGI/Cray*

10:00 Advancing with UNICOS Toward IRIX
Barry Sharp, BCS

Rhein (4.3.11)

3B Performance and Evaluation

Chair: Jeffery A. Kuehn, NCAR

9:00 Parallel Job Performance in a Time Sharing Environment
Dave McWilliams, UIUCNCSA

9:30 Industrial Applications on MPP Architectures
Jörg Stadler, debis Systemhaus GmbH

10:00 Performance Co-Pilot and Large Systems Performance
Ken McDonell, SGI/Cray

Monrepos (4.4.25)

3C Operating Heterogeneous Sites

Chair: Mike Brown, EPCC

9:00 SGI/Cray Monitoring Tools
Randy Lambertus, SGI/Cray

9:30 Industry Directions in Storage
Mike Anderson, SGI/Cray

10:00 (cont.)

10:30–11:00 Break

Schiller-Saal

4 General Session
Chair: Eric Greenwade, INEEL

11:00 The New Silicon Graphics
Rick Belluzzo, CEO, SGI

12:00 Questions and Answers
Beau Vrolyk, SGI

12:30–2:00 Lunch

Hohenzollern (4.4.22)

5A IRIX: Origin Software Road Maps

Chair: Nicholas Cardo, SS-SSD

2:00 IRIX on High End Origins: Plans & Status
Jim Harrell, SGI/Cray
2:30 MPT and Cluster Software Update
Karl Feind, SGI/Cray
3:00 Update of System Management Software for Large Origin Systems
Daryl Coulthart, SGI/Cray

Rhein (4.3.11)

5B Exploiting Parallelism

Chair: Larry Eversole, JPL

2:00 OpenMP Programming Model
Ramesh Menon, SGI/Cray
2:30 OpenMP: a Multitasking and Autotasking Perspective, *Neal Gaarder, SGI/Cray*
3:00 Expressing Fine-Grained Parallelism Using Fortran Bindings to Posix Threads, *Henry Gabb, ARMY/WES*

Schiller-Saal

5C User Services

Chair: Denise Brown, USARL

2:00 Recycling Instructor-led Training on the WWW, *Leslie Southern, OSC*
2:30 Design Your Own Webzine: A Practical Guide, *Lynda Lester, NCAR*
3:00 (cont.)

3:30–4:00 Break

Hohenzollern (4.4.22)

6A GigaRings, Fill-ups & Tune-ups

Chair: Hartmut Fichtel, DKRZ

4:00 GigaRing Disk Update
Tom Hotle, SGI/Cray
4:30 Performance Tips for GigaRing Disk IO
Kent Koeninger, SGI/Cray
5:00 Cray Data Migration Facility & Tape Management Facility Update
Neil Bannister, SGI/Cray

Rhein (4.3.11)

6B Programming Environments and Models

Chair: Hans-Hermann Frese, ZIB

4:00 Programming Environment Status and Update, *Sylvia Crain, SGI/Cray*
4:30 IRIX Programming Tools Update
Bill Cullen, SGI/Cray
5:00 Cray Product Installation and Configuration, *Scott Grabow, SGI/Cray*

Monrepos (4.4.25)

6C SGI/Cray Operations Tools and Storage Technologies

Chair: Mike Brown, EPCC

4:00 DCE/DFS: An Environment for Secure & Uniform Access to Data and Resources
Dieter Mack, RUS
4:30 Monitoring at NCSA
Tom Roney, UIUCNCSA
5:00 A Common Operating Environment
Bill Middlecamp, SGI/Cray

Hohenzollern (4.4.22)

7A   Networking the Engines

Chair: Hans Mandt, BCS-ASL

9:00   Scalability and Performance of Distributed I/O on Massively Parallel Processors, *Peter W. Haas, RUS*

9:30   The UNICORE Project: Uniform Access to Supercomputing over the Web
Jim Almond, ECMWF
Mathilde Romberg, KFA

10:00   Cray Networking Update
Michael Langer, SGI/Cray

Rhein (4.3.11)

7B   Parallel and Distributed Computing

Chair: Hans-Hermann Frese, ZIB

9:00   MPI Regression Testing and Migration to Origin 2000, *Terry Nelson, SS-SSD*

9:30   Distributed Supercomputing Services in a Heterogeneous Environment
Peter Morreale, NCAR

10:00   High Performance Fortran for the T3E and Origin
Douglas Miles, Portland Group

Schiller-Saal

7C   Visualization

Chair: L. Eric Greenwade, INEEL

9:00   Visualization of Astrophysical Data with AVS/Express, *Jean Favre, ETHZ*

9:30   Interactive Direct Volume Rendering on the O2K, *John Clyne, NCAR*

10:00   VRML for Visualization
Jim Johnson, SGI/Cray

10:30–11:00   Break

Schiller-Saal

8   General Session
Chair: Barbara Horner-Miller, ARSC

11:00   CUG Elections
Jean Shuler, LLNL

11:40   Keynote Address—25 Years of Computer Aided Engineering at Daimler Benz in Stuttgart, Germany
Michael Heib, Manager HWW

12:25   CUG Election Results

12:30–2:00   Lunch

Hohenzollern (4.4.22)

9A — High Speed Data!

Chair: Harmut Fichtel, DKRZ

2:00   High Performance Network Backup &
Restore with HYPERtape
Uwe Olias, UNIKIEL

2:30   Functional Comparison of DMF and
HPSS, *Gerhard Rentschler, RUS*

3:00   The Rebirth of DMF on IRIX
Alan Powers, NAS

Rhein (4.3.11)

9B — Powering Science

Chair: Larry Ebersole, JPL

2:00   Material Science on Parallel Computers
(T3E/900), *Andrew Canning, NERSC*

2:30   Clustering T3Es for Metacomputing
Applications, *Michael Resch, RUS*

3:00   Synchronization Using Cray T3E Virtual
Shared Memory

Monrepos (4.4.25)

9C — SGI/Cray Service Plans and Q & A Panel

Chair: Dan Drobnis, SDSC

2:00   SGI/Cray Service Report
Bob Brooks, SGI/Cray

2:30   SGI/Cray Q&A Panel
Moderator: Charlie Clark, SGI/Cray

3:00   (cont.)

3:30–4:00   Break

Schiller-Saal

10A — IRIX: The Long Road Ahead

Chair: Nicholas Cardo, SS-SSD

4:00   Cellular IRIX: Plans & Status
Gabriel Broner, SGI/Cray

4:30   Year 2000 Compliance for UNICOS &
UNICOS/mk
Dennis Arason and Bruce Schneider, SGI/Cray

5:00   (cont.)

10B — Performance and Evaluation

Chair: Michael Resch, RUS

4:00   High-Performance I/O on Cray T3E
Ulrich Detert, KFA

4:30   Cray T90 versus Tera MTA
Jay Boisseau, SDSC

5:00   (cont.)

10C — Joint Session— Software Tools and User Services

Chair: Hans-Hermann Fresse

4:00   Supercomputing—Strategies and Applications at the High-Performance Computing Center, Stuttgart
Alfred Geiger, RUS

4:30   Customer Scenarios for Large Multiprocessor Environments
Fritz Ferstl, Genias Software GmbH

5:00   How CUG Members Support Their Users
Barbara Horner-Miller, ARSC

CUG Night Out, 7:00 p.m.
Busses depart from in front of the Conference Center at 6:30
(Remember to bring your invitation letter from debis!)

Hohenzollern (4.4.22)

IIA The Power of the Engine: Users and the O/S

Chair: Ingeborg Weidl, MPG

9:00 The Age-Old Question of How to Balance Batch and Interactive
Barry Sharp, BCS

9:30 MISER: User Level Job Scheduler
Kostadis Roussos, SGI/Cray

10:00 Serving a Demanding Client While Short on Resources, *Manfred Stolle, ZIB*

Rhein (4.3.11)

IIB Applications and Algorithms

Chair: Larry Ebersole, JPL

9:00 Application Roadmap
Jef Dawson, SGI/Cray

9:30 Mathematical Methods for Mining in Massive Data Sets
Helene E. Kulsrud, CCR-P/IDA

10:00 (cont.)

Monrepos (4.4.25)

IIC GigaRing Operations

Chair: Mike Brown, EPCC

9:00 GigaRing Configuring
Michael Langer, SGI/Cray

9:30 GigaRing Systems Monitoring
Birgit Naun, Thomas Plaga, KFA; Ralph Krotz, SGI/Cray

10:00 Open

10:30–11:00   Break

Schiller-Saal

I2 General Session
Chair: Bonnie Hall, LANL

11:00 SGI/Cray Service Report
Ken Coleman, SGI/Cray

11:30 SGI/Cray Joint Software Report
Mike Booth, SGI, Denice Gibson, SGI/Cray

12:30–2:00   Lunch

Hohenzollern (4.4.22)

I3A — Hauling Big Data

Chair: Hartmut Fichtel, DKRZ

2:00 Mass Storage at the NCSA: DMF and Convex UniTree
Michelle Butler, NCSA

2:30 Towards Petabytes of High Performance Storage at Los Alamos
Gary Lee, LANL

3:00 Storage and Data Management—Big Data Solutions
Ken Hibbard, SGI/Cray

Rhein (4.3.11)

I3B — Secure & Controlled Computing Engines

Chair: Bonnie Hall, LANL

2:00 Securing the User's Work Environment
Nick Cardo, SS-SSD

2:30 The State of Security for UNICOS & IRIX, *Jay McCauley, SGI/Cray*

3:00 IRIX Accounting Limits and UDB Functionality
Jay McCauley, SGI/Cray

Monrepos (4.4.25)

I3C — User Services

Chair: Leslie Southern, OSC

2:00 NPACI User Services
Jay Boisseau, SDSC

2:30 Cray Training Update
Bill Mannel, SGI/Cray

3:00 Pushing/Dragging Users Towards Better Utilization
Guy Robinson, ARSC

3:30–4:00 Break

Hohenzollern (4.4.22)

I4A — Open Meetings

Rhein (4.3.11)

I4B — Open Meetings

Monrepos (4.4.25)

I4C — Open Meetings

4:00–5:30 Please plan to attend the **Open Meetings** to discuss topics of interest to each Special Interest Group.

6:00 Also plan to attend the **Program Steering Committee Meeting** and help us prepare the program for our next conference.

The locations for these meetings will be posted on the bulletin board.

Hohenzollern (4.4.22)

15A IRIX: From Plans to Reality

Chair: Nick Cardo, SS-SSD

9:00 Getting It All Together
Daryl Grunau, LANL

9:30 Integrating an Origin2000 into a Cray Data Center
Chuck Keagle, BCS

Rhein (4.3.11)

15B Applications and Algorithms

Chair: Richard Shaginaw, BMSPRI

9:00 XVM–Extended Volume Management
Colin Ngam, SGI/Cray

9:30 The Good, the Bad, and the Ugly Aspects of Installing New OS Releases
Barry Sharp, BCS

10:00 Origin Craylink Partitioning
Steve Whitney, SGI/Cray

Schiller-Saal

15C Visualization

Chair: L. Eric Greenwade, INEEL

9:00 CUG Video Theater

9:30 (cont.)

10:00 Visualization of 3 Dimensional Material Science Applications
L. Eric Greenwade, INEEL

10:30–11:00   Break

Schiller-Saal

16 General Session

Chair: CUG Vice-President

11:00 CUG SIG Reports: *CUG Vice-President*

11:10 Parallel and Distributed Development and Simulation of Atmospheric Models
V. Mastrangelo and I. Mehilli, CNAM-Université;
F. Schmidt, M. Weigele, J. Kaltenbach, A. Grohmann, and
R. Kopetzky, University of Stuttgart

11:40 SGI/Cray Hardware Report and Hardware Futures
Steve Oberlin, SGI/Cray

12:40 CUG Next Steps: CUG 99 in Minneapolis, MN
John Sell, MSC

12:50 Closing Remarks: *CUG President*

1:00 *End*

Abstracts

Tutorial I

8:30

UNICOS/IRIX Differences in DMF

Paul Ernst, SIG/Cray

This tutorial presents the differences between the UNICOS and IRIX implementations of the Data Migration Facility. Topics covered will include feature differences, installation and configuration changes, DMF tape interface information (TMF and OpenVault), and conversion assistance for sites considering changing platforms from UNICOS to IRIX.

Tutorial II

8:30

Scheduling and Configuration for Origin 2000

Jim Harrell, SGI/Cray

This talk will provide an understanding of large (64 processor and above) Origin configuration and tuning. The basics of Origin configuration and tuning will be explained. SGI/Cray's understanding and experience running large Origins will be used to provide guidance for customers on these large systems.

Tutorial III

8:30

Scheduling and Configuration Tuning for the T3E

Jim Grindle, SGI/Cray

The T3E, while similar to the T3D for application programming, offers a host of new possibilities for system configuration. This tutorial will address these possibilities and present knowledge obtained from recent experience. The configuration includes such issues as distribution of the OS servers, the placement of OS servers on the OS PEs, and the organization of pcache. A key question, the use of political scheduling, will be described. A walk-through of how to make key configuration choices and a discussion of their effects will be presented. Recommendations for optimal settings will be made. Since our knowledge is incomplete, input or feedback from sites about their experiences will be welcomed as part of the tutorial.

Tutorial IV

8:30

Visualization from Theory to Practice

L. Eric Greenwade, INEEL

The aim of this tutorial is to provide a greater understanding of the process and products of visualization. After a brief overview of the what, why and when of visualization, a detailed example will be presented. Starting from the problem statement, gaining new insight in a complex data set, each step of the process will be detailed with the pitfalls and intermediate stages described. The end result will then be analyzed as to the effectiveness of the output.

Tutorial V

11:00

Contemporary Issues in Network Security

Jay McCauley, SGI/Cray

This tutorial reviews the current state of the art in security for networked computer systems. The material covered will include: an introduction to firewalls and proxies, internet authentication technologies, encryption based solutions including SSL and ssh, and security screening and monitoring tools. Some of the new security features found in the IRIX™ 6.5 Operating System and related products will be discussed including capabilities, access control lists, and the new System Security Scanner found in WebForce™ product. While no formal text is required, Bellovin and Cheswick's "Firewalls and Internet Security: Repelling the Wily Hacker" (new edition soon to be published) is an excellent overview of some the topics covered.

Monday June 15, 1998

Tutorial VI

11:00

Performance Optimization for the Origin 2000

Jeff Brooks, SGI/Cray

This tutorial will describe the Origin 2000 architecture and the resulting programming concepts. Topics include: Features of the Origin 2000 Hardware, MIPS f90 compiler and its command line options, performance tools, single processor tuning and multi-processor tuning. The tutorial assumes "Cray-centric" programming expertise.

Tutorial VII

11:00

Mathematical Libraries

Bill Harrod, SGI/Cray

This tutorial will discuss a variety of scientific libraries routines for solving linear algebra and signal processing problems. The LAPACK package and SGI/Cray FFT routines will be highlighted. LAPACK provides a choice of algorithms mainly for dense matrix problems that are efficient and portable on a variety of high performance computers. Examples will be provided for converting LINPACK or EIS-PACK subroutine usage to the appropriate LAPACK subroutine. We will also discuss other migration issues such as porting LIBSCI based application codes to the new Origin 2000 scientific library called SCSL. The tutorial will also address the challenge facing designers of mathematical software in view of the development of highly parallel computer systems. We shall discuss ScaLAPACK, a project to develop and provide high performance scaleable algorithms suitable for highly parallel computers.

OPEN MEETINGS

4:00, 4:30, and 5:00

Please plan to attend the Open Meetings to discuss topics of interest to each Special Interest Group. Check the conference bulletin board for details.
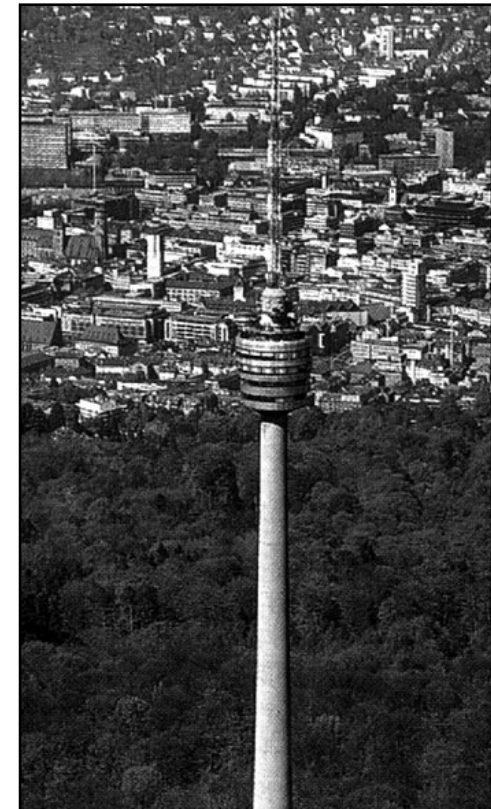
Abstracts

3A The Engine's Core: Operating Systems

Chair: Ingeborg Weidl, MPG

9:00

PVP UNICOS Status and Update

Patti Langer, SGI/Cray

This talk will cover current release and support plans, reliability/stability efforts, future plans as well as T90P and J90++ updates.

9:30

T3E and UNICOS/mk Status and Update

Jim Grindle, SGI/Cray

This talk will cover current status of system issues and their progress. We will specifically cover scheduling, GRM, and resiliency issues in UNICOS/mk on the T3E, along with a project update with release and support plans.

10:00

Advancing with UNICOS Toward IRIX

Barry Sharp, BCS

During the past seven years of running UNICOS at Boeing on X-MP, YMP, and T90 systems many enhancements to the OS and System Admin, and User-Level facilities have been made. This paper will outline the rationale for these and provide brief descriptions on their implementation. The intent is to share the inventory with other UNICOS member sites and discuss those that might be applicable or of interest to the Origin 2000 IRIX system.

3B Performance and Evaluation

Chair: Jeffery A. Kuehn, NCAR

9:00

Performance of Thread-Parallel and Message-Passing Jobs in a Time-Sharing Environment on Origin 2000 Systems

Dave McWilliams, UIUCNCSA

At NCSA, we discovered that thread-parallel, gang-scheduled Origin 2000 jobs can use as much as 5 times as much CPU time when the load average on a system is high (e.g., 96 on a 64-processor system). We have been working with SGI as they make changes in the IRIX process scheduler to fix the problem. Although the work is still in progress, we have already seen improvements in job performance. We are planning to explore the performance of MPI jobs on a heavily loaded system in the near future. We will also discuss efforts to limit the load average on IRIX systems.

9:30

Cray T90 versus Tera MTA: The Old Champ Faces a New Challenger

Jay Boisseau, SDSC

The T90 represents the latest in CRI's line of parallel-vector supercomputers and still reigns as the champion in terms of performance on many production codes (for which parallel versions have not and perhaps will not be developed). Tera has recently delivered their first Multi-Threaded Architecture (MTA) to SDSC and hopes to provide a truly scalable shared-memory parallel platform that will compete favorably in performance and ease of programming against the T90. We present the results of running a few efficient T90 user codes on the MTA without and with optimization and compare to their T90 performance, with initial analysis of the reasons for performance differences.

10:00

Performance Co-Pilot and Large Systems Performance

Ken McDonell, SGI/Cray

Brief Overview of PCP features, scope of performance issues it is trying to address, canned demos, availability. Addressing Large System Performance.

- getting performance data out of the application and into PCB
- libpcp_trace
- PMDA construction and libpcp_pmda
- customization (leverage the PCP building blocks to enhance the tools for performance management of your application)
- pmchart, pmview,mpgadgets
- pmlogger
- pmie
- scaleable performance visualizations
- extensibility
- correlating resource demands across the hardware, operating system, service and application layers
- measuring quality of service
- Case Study—building an MPI Performance Monitoring Toolkit from PCP parts

3C Operating Heterogeneous Sites

Chair: Mike Brown, EPCC

9:00

J90, T90, C90, T3, GigaRing, Origin2000 Monitoring Tools

Randy Lambertus, SGI/Cray

The presentation includes product description, application implementation, system monitoring, problem notification, service response actions and real-life examples, and concentrates on the present products Watchstream, Watchlog and Availmon and the related interface for each application sub-function.

Ideas for future development efforts, such as site system metric gathering and expanded notification options, will be opened for general discussion with the audience.

9:30

A Look at Disk and Tape Futures from the SGI/Cray Perspective

Mike Anderson, SGI/Cray

The disk drive industry continues to undergo changes at an ever increasing pace. Competition and technology advances continue to change the products produced and shorten the product life cycles. The tape drive industry is also undergoing big changes. An overview of these changes and the impact on SGI/Cray disk and tape products will be presented

STUTTGART

5A IRIX: Origin Software Road Maps

Chair: Nicholas Cardo, SS-SSD

2:00

IRIX on High End Origins: Plans & Status

Jim Harrell, SGI/Cray

This is an update of the plans and status for IRIX OS support of the high end systems. The primary focus will be the status of current system software releases and the plans for future releases and upgrades. The status and scalability of IRIX on more than 64 processors will be discussed.

2:30

MPT and Cluster Software Update for IRIX, UNICOS/mk, and UNICOS

Karl Feind, SGI/Cray

IRIX cluster software continues evolving to add both high availability and scalability functions. This talk briefly summarizes recent or planned changes with the FailSafe, NQE, IRISconsole, Array Services, and MPT products. The majority of the talk centers on MPI and SHMEM in MPT 1.2. In addition to review of new features, performance of MPI on Origin and SHMEM on Origin and T3E will be overviewed.

Abstracts

3:00

Update of System Management Software for Large Origin Systems

Dan Higgins, SGI/Cray

Support for large Origin systems has dramatically improved in the last year. This paper describes the use of IRIX systems software to improve management of large Origin systems. The software is IRIX 6.4, MPI, Array services, checkpoint, NQE jobs limits, MISER and AXMAN for interactive limits. There will be additional improvements in 1998. UNICOS jobs limits, the UDB, job notion and extended accounting will be added to IRIX.

5B Exploiting Parallelism

Chair: Richard Shaginaw, BMSPRI

2:00

OpenMP Programming Model

Ramesh Menon, SGI/Cray

OpenMP is an application program interface (API) for shared-memory parallel programming. Pioneered by SGI, it is fast becoming a de facto industry standard, as evidenced by the large number of hardware and software vendors endorsing the standard. The functionality is designed to enable programmers to write coarse grain, scaleable, shared-memory parallel programs. This talk will present the why, what, and how of OpenMP.

2:30

OpenMP: a Multitasking and Autotasking Perspective

Neal Gaarder, SGI/Cray

The new OpenMP standard for shared-memory parallelism is a collection of standard compiler directives, library routines, and environment variables for shared memory parallelism. This paper discusses OpenMP and compares it with current PVP autotasking and multitasking capabilities.

3:00

Expressing Fine-Grained Parallelism Using Fortran Bindings to Posix Threads

Henry Gabb, ARMY/WES

In dynamics simulations, the through-space interactions between particles, which must be calculated every time step, consume the bulk of computational time. These calculations typically occur in a single, large loop containing many data dependencies. However, the iterations are often independent, so fine-grained parallelism should confer a significant performance gain. Pthreads have significant advantages over compiler directives, which often create separate UNIX processes. Multiple threads exist in a single process and require less system overhead. Also, threads are not linked to physical processors as is often the case for compiler directives. Multiple threads residing on a single processor give better resource utilization (e.g., separate threads doing computation and I/O operations). Performance and programming issues that arise when expressing fine-grained parallelism on SGI SMP and Cray T3E architectures will be discussed.

5C User Services

Chair: Denice Brown, USARL

2:00

Recycling Instructor-Led Training on the WWW

Leslie Southern, OSC

Web-based formats have great potential for broadening the reach of traditional instructor-based training courses. To enhance training materials on the WWW, the Ohio Supercomputer Center is experimenting with audio and slide combinations on the WWW. The objective is to re-use these courses to reach a larger audience at any time and any place. In addition to case studies, descriptions of experiences, hardware, software, and tools will be included. Evaluation methods and available user feedback will also be presented.

Abstracts

2:30

Design Your Own Webzine: A Practical Guide

Lynda Lester, NCAR

This paper gives concrete examples of what to do and what not to do in creating a newsletter (Webzine) for users on the WWW. We discuss navigation aids; methods of information retrieval (search engine, index, contents); principles of typography and page design; use of graphics and animations; cross-platform "gotchas"; writing for the Web; review cycles; URLs and file structure; and user notification procedures. What does "interactive" mean? Do readers still want hardcopy? How far will people click? We will cover all these points and more in an entertaining and pragmatic talk on the basic concepts of Web design.

6A GigaRings, Fill-ups & Tune-ups

Chair: Hartmut Fichtel, DKRZ

4:00

GigaRing Mass Storage Update

Michael Langer, SGI/Cray

This talk will focus on the GigaRing systems Mass storage plans, including release plans, support plans, stability and performance.

4:30

Performance Tips for GigaRing Disk IO

Kent Koeninger, SGI/Cray

This talk will present techniques on CRAY T90, CRAY J90 and CRAY T3E systems to maximize this GigaRing disk performance, including FCN RAID disks, pcache on T3Es, SSD-T90 caching on T90s, memory caching on J90 systems user-memory caching (FFIO), disk configuration tips, and other UNICOS features for fast IO. The recent FCN RAID performance improvements enabled measurements and characterization of these important I/O techniques.

5:00

Cray Data Migration Facility & Tape Management Facility Update

Neil Bannister, SGI/Cray

This paper will review status and plans for both Cray and SGI platforms. Since the last CUG meeting DMF was released. This has presented the DMF team with new challenges which will be covered in this paper. The paper will also report progress and release plans for the TMF and CRL products.

6B Programming Environments and Models

Chair: Hans-Hermann Frese, ZIB

4:00

Programming Environment Status and Update

Sylvia Crain, SGI/Cray

This talk will cover current release and support plans for the Programming Environment for the MPP and PVP platforms. It will also discuss future plans in this area including migration efforts.

4:30

IRIX Programming Tools Update

Bill Cullen, SGI/Cray

In November 1997, the IRIX programming tools changed direction. This talk will cover the current status and road map for Workshop, Speedshop, MPF, RapidApp and CosmoCode.

5:00

Cray Product Installation and Configuration

Scott Grabow, SGI/Cray

The intent of this talk is to provide insight into how the Common Installation Tool (CIT) is being used to make installation of software follow a similar process, and address common problems that have been seen in the migration to CIT and CD-ROMs for media distribution. In addition, examples of how these process changes affected various manuals will be shown, and how the goal is to follow a task oriented approach, and to separate installation and system configuration issues into different manuals.

6C Cray/SGI Operations Tools and Storage Technologies

Chair: Mike Brown, EPCC

4:00

Integration of Origin 2000s, T90s & J90's at USARL

Albert G. Edwards IV, USARL

The ARL MSRC has an environment where the Origin 2000, T90 and the J90 have been integrated into a common computing environment. This experience has not been without its challenges. The ARL MSRC has had to migrate many users and codes from an 8 node Power Challenge Array and a Cray-2 to the Origin 2000, T916, J932 and the J916. The J916 is functioning as the fileserver for the Origins and T90's. This talk will focus on the experiences and lessons learned in the installation and integration of these computing platforms.

4:30

Toward an Integrated Monitoring Scheme

Tom Roney, UIUCNCSA

The National Center for Supercomputing Applications (NCSA) employs a virtual operator (Voper) on UNIX operating systems. Voper detects potential system and application problems, and displays warning messages for administrative attention. Voper is being further developed for interactive use, to provide automated implementation of counter measures against reported problems. Voper and all other tools used to monitor NCSA systems are being collected to run under a single software package, the Computer Associates' Unicenter. NCSA effectively pilots a production environment, and is progressing toward an integrated monitoring scheme.

5:00

SGI/Cray Plans for Operations Commonality across MPP, Vector, and Scaleable Node Environments

Bill Middlecamp, SGI/Cray

Cray plans to offer a Common Operating Environment across traditional Cray MPP, parallel vector, and scaleable node product families. The Common Operating Environment, along with the Common Supercomputing API, are intended to provide interoperability and ease customer transition among the product families. A Common Operating Environment API specification will be available to customers at this CUG.

Abstracts

7A The Network & Computing Engines

Chair: Hans Mandt, BCS-ASL

9:00

Scalability and Performance of Distributed I/O on Massively Parallel Processors

Peter W. Haas, RUS

Supercomputer network configurations, both internal and external, have developed over several orders of magnitude in performance in recent years. At the same time, the scope of networks has been broadened to accept not only the traditional vector supercomputer but also massively parallel systems of various types, file servers, workstation clusters, visualization laboratories and multimedia technology.

Networking and file system I/O on MPPs have been confined to a limited number of system level processors in the past leading to well known bottlenecks, especially with the execution of network protocols. There are elegant ways, however, to co-locate server processes on user nodes, which will enable truly distributed I/O.

The High Performance Storage System (HPSS) testbed at RUS is taken as an example to illustrate the basic principles. This testbed covers all major computing platforms.

9:30

The UNICORE Project: Uniform Access to Supercomputing over the Web

Jim Almond, ECMWF
Mathilde Romberg, KFA

Supercomputers are becoming more powerful, but also more centralized in fewer centers. To fully utilize the potential of such facilities, more uniform, secure, and user friendly access via the Internet is needed. In addition to these generalities, this talk will describe the UNICORE project, a large collaboration dedicated to the implementation of a prototype addressing the above goals.

10:00

Cray Networking Update

Michael Langer, SGI/Cray

This talk will cover future Cray/SGI networking plans, including release plans, support plans, stability and performance.

7B Parallel and Distributed Computing

Chair: Hans-Hermann Frese, ZIB

9:00

MPI Regression Testing and Migration to Origin 2000

Terry Nelson, SS-SSD

The computing industry, for economic and technical reasons, is moving inexorably towards an increasingly parallel environment. One of the major paradigms to accomplish this is message passing, and one of the major tools in this area is MPI. This paper will describe a set of Fortran and C regression tests which test MPI functionality and performance. A series of runs on a C90, J90, and Origin 2000, as well as tests with production sized jobs, will be described. Particular attention will be given to the migration of programs from the C90 and J90 environments to the Origins.

9:30

Distributed Supercomputing Services in a Heterogeneous Environment

Peter Morreale, NCAR

The NCAR Distributed Computing Services (DCS) project is a five year effort aimed at providing NCAR supercomputing services on

local users desktops. The initial effort has focused on providing NCAR Mass Storage System (MSS) services to major compute servers including Cray, SGI, IBM, and Sun systems. The DCS system is designed around OSF's DCE software. DCS makes liberal use of the DCE Remote Procedure Call (RPC) mechanism, as well as the Cell Directory Service (CDS). This paper discusses the design of the NCAR DCS system as currently implemented as well as future directions.

10:00

High Performance Fortran for the T3E and Origin

Douglas Miles, Portland Group

No abstract available

7C Visualization

Chair: L. Eric Greenwade, INEEL

9:00

Visualization of Astrophysical Data with AVS/Express

Jean Favre, ETHZ

In the frame of astrophysical applications Euler equations including source terms are solved numerically in two and three dimensions. The data are computed on a J90-cluster by an adap-

tive mesh code with a high degree of vectorization and parallelization. The output data are stored in a multi-level hierarchy with solutions at different levels of spatial resolution.

These time-dependent simulations impose very high visualization constraints. We use AVS/Express to integrate custom developments required for memory, CPU, graphics resources and remote data access imposed by the application. We present the implementation of the visualization environment on SGI workstations with our 512-Tbyte storage facility.

9:30

Interactive Direct Volume Rendering of Time-Varying Data on the Origin 2000

John Clyne, NCAR

Direct Volume Rendering (DVR) is a powerful volume visualization technique for exploring complex three and four dimensional scalar data sets. Unlike traditional surface fitting approaches to volume visualization, which map volume data into geometric primitives and can benefit greatly from widely-available commercial graphics hardware, computationally-expensive DVR is performed, with rare exception, exclusively on the CPU. Fortunately DVR algorithms tend to parallelize readily. Much prior work has been done to produce parallel volume renderers capable of visualiz-

ing static data sets in real-time. Our contribution to the field is the development of parallel software that takes advantage of high-bandwidth networking and storage to deliver volume rendering of time-varying data sets at interactive rates. We discuss our experiences with the software on the Origin 2000 class of supercomputers.

10:00

VRML for Visualization

Jim Johnson, SGI/Cray

VRML, the Virtual Reality Modeling Language, is heading for a browser near you. VRML promises a write once, use everywhere, capability for visualizing the results of engineering and scientific calculations. The same results may be visualized on a multitude of platforms, locally or over the web, using familiar web browsers. The format and functionality are unchanged, only the performance and capacity vary. Simulation results can be packaged in a VRML format file to be loaded by a browser. Alternatively, a Java applet can read an existing data format and inject the data into a running VRML-capable browser.

Abstracts

9A High Speed Data!

Chair: Hartmut Fichtel, DKRZ

2:00

High Performance Network Backup & Restore with HYPERtape

Aindrias Wall, MultiStream Systems Inc.

The fastest backup & restore for heterogeneous platforms, including Cray UNICOS and SGI IRIX as server platforms. It takes full advantage of their outstanding I/O performance. Overall, HYPERtape covers the most various client platforms and is highly scaleable.

2:30

Functional Comparison of DMF and HPSS

Gerhard Rentschler, RUS

RUS is looking for a new data management solution for its supercomputing area. Different solutions are being evaluated. Among them are HPSS and DMF for IRIX. This talk will compare both products and point out the strengths and weaknesses.

Experiences in a real life environment are also presented.

3:00

The Rebirth of DMF on IRIX

Alan Powers, NAS

The Numerical Aerodynamic Simulation Facility (NAS) at NASA Ames Research Center (AMES) is in the process of testing DMF 2.6 on a Power Challenge and an Origin 2000. These systems are connected to a STK 4400 silo using SCSI STK 9490 tape drives. A list of new features and differences will be covered. A simple benchmark will be used to compare the performance of IRIX DMF and UNICOS DMF.

9B Powering Science

Chair: Richard Shaginaw, BMSPRI

2:00

Material Science on Parallel Computers (T3E/900)

Andrew Canning, NERSC

Three of the most heavily used quantum methods in material science (Plane Wave, Density Functional Theory, and the Tight Binding approach) will be discussed, along with their implementation on a variety of different parallel machines, including the 544-processor Cray T3E at NERSC. The new parallel algorithms developed to run these codes efficiently on parallel machines will be presented. Applications to large problems in material science will also be presented.

2:30

Clustering T3Es for Metacomputing Applications

Michael Resch, HLRS

Having developed a library that allows running MPI applications on a cluster of MPPs, we have done some experiments with several applications. So far we have good results for a CFD application; but even more promising are the results for Monte Carlo simulations.

3:00

Synchronization Using Cray T3E Virtual Shared Memory

Miltos Grammatikakis, KFA

We consider mutual exclusion on the Cray T3E shared memory using various atomic operations and algorithms. Our current implementations, when compared to the Cray shmem_lock functions, indicate time improvements of at least 2 orders of magnitude for a 64-processor T3E/900. Our comparisons are based on both synthetic and actual code for concurrent priority queues. Software barrier performance is also briefly examined.

Abstracts

9C SGI/Cray Service Plans and Q & A Panel

Chair: Dan Drobnis, SDSC

2:00

SGI/Cray World-Wide Customer Service Plans

Bob Brooks, SGI/Cray

World-Wide Service at SGI has undergone many changes since the merger with Cray. This talk will describe those changes, provide a description of current service delivery strategies, and define plans and initiatives that will affect the delivery of service in the future.

2:30

SGI/Cray Q&A Panel

Moderator: Charlie Clark, SGI/Cray

Representatives of SGI/Cray Service, Development, and Support organizations will discuss issues and questions raised by the CUG Survey and at the OpsSIG Business Meeting, and field questions and comments from the floor. This is a traditional wrap-up to OpsSIG sessions, and often a lively dialog.

10A IRIX: The Long Road Ahead

Chair: Nicholas Cardo, SS-SSD

4:00

Cellular IRIX: Plans & Status

Gabriel Broner, SGI/Cray

Future operating systems developed by Cray and SGI for the supercomputer and server spaces will be based on Cellular IRIX technology. Cellular IRIX is an evolution of IRIX which incorporates supercomputer features existing on UNICOS and UNICOS/mk. Key features of Cellular IRIX are its increased fault containment and its scalability to thousands of processors. This talk will cover the architecture of Cellular IRIX and the contents and status of the first operating system releases using this technology.

4:30

Year 2000 Compliance for UNICOS & UNICOS/mk

Dennis Arason, SGI/Cray

This talk will provide information as to what testing was done to ensure 2000 compliancy for UNICOS and U/mk systems. It will cover the platforms tested and what testing was done. It will also cover the products tested and open issues.

10B Performance and Evaluation

Chair: Michael Resch, RUS

4:00

Experiences with High-Performance I/O on Cray T3E

Ulrich Detert, KFA

No abstract available

4:30

Comparing T90 and T3E Performance on Optimized Fortran 90 and HPF Codes

Jay Boisseau, SDSC

The National Partnership for Advanced Computational Infrastructure (NPACI) provides access to a CRAY T90 and two CRAY T3Es (and other parallel computing platforms). As the reasons (and pressures) to migrate codes to parallel platforms increase, some users are interested in converting production Fortran codes to HPF to minimize development time. We compare single- and multi-CPU performance results for several optimized T90 codes to their HPF versions optimized for the T3E to illustrate the performance one might expect in migrating to HPF. We also provide information on tuning HPF codes for maximum performance on the CRAY T3E.

5:00

Industrial Applications on MPP Architectures

Jörg Stadler, debis Systemhaus GmbH

HWW (Höchstleistungsrechenzentrum für Wissenschaft und Wirtschaft GmbH) is a large German supercomputing facility, that offers its computer resources to scientific as well as industrial users. This talk reports on the applications that industrial users run on the HWW equipment. An increasing demand for computer resources from commercial customers eventually led to the installation of an additional machine in March 1998. All of the demand is still exclusively met by vector machines, although HWW offers two large MPPs (Cray T3E and IBM SP2). These MPPs are still being evaluated by the industrial users. On the one hand, the users are interested in the MPPs, because they know that MPPs offer advantages in terms of price/performance ratio, raw computer power, or main memory size. In addition, the main applications are available on MPPs, so testing and evaluation is straightforward. On the other hand, the MPP technology created it's own breed of problems. Consider for example crash simulation; although this is a well known industrial supercomputer application, it still is the area with the largest demand for computer resources. Naturally, several efforts have been made to run crash simulations on

MPPs. These efforts have shown that just porting an application to a new architecture is sometimes not enough. More specifically, the numerical sensitivity of crash models led to several new difficulties. For example, dynamic load balancing, normally considered good programming practice on MPPs, gives rise to results that vary significantly from run to run. Furthermore, the results depend on the number of processors used. CFD users are also evaluating the use of MPPs. We have run some successful benchmarks and demonstrations (including a model with an 8GB dataset) in that area. As a summary, it can be stated that industrial users are having a close and concentrated look at MPPs right now, but they are not running their production calculations on these machines at this time.

10C Joint Session— Software Tools and User Services

Chair: Hans-Hermann Frese, ZIB

4:00

Supercomputing—Strategies and Applications at the High-Performance Computing Center, Stuttgart

Alfred Geiger, HLRS

High Performance Computing in Germany is undergoing a major transition from tens of

middle-class centers to mainly four high-end centers. The center in Stuttgart, HLRS, is oriented towards engineering-applications. It was therefore a natural step to enable synergetic effects in acquisition, operation and usage by outsourcing the supercomputers into a joint-venture with industrial partners.

The main application-fields of the center in Stuttgart are CFD, reacting flows, structural mechanics and electromagnetics. In most of these applications it could be proven, that large-scale engineering-simulations can scale very well to hundreds of nodes on scalar as well as on vector-based parallel systems. Examples of such simulations are:
- Reentry of space-vehicles
- Combustion-simulation in coal-fired power-plants
- Crashworthiness-simulations
- Radar cross-section
- Internal combustion-engines

For industrial users an optimal embedding of HPC into the development-process is necessary. Most companies have development-groups all over the world, making tools for cooperative working in HPC environments unavoidable if time to market is a critical issue. Examples of EC-funded projects in this area are presented.

There are problems, that cannot be simulated on even the largest supercomputers in the world. Metacomputing-scenarios, though not

Abstracts

very efficient, are the only way in such cases. The results and techniques of such projects will be explained.

4:30
Customer Scenarios for Large Multiprocessor Environments

Fritz Ferstl, Genias Software GmbH

Genias Global Resource Director (GRD) was developed to meet the demands of large multiprocessor sites with several 100s of CPUs and users. At the Army Research Laboratory (ARL), GRD is installed on several types of Cray and SGI machines.

The talk will describe how GRD
- integrated these machines into one environment, and enabled administrators to define global management policies. Policy decisions for individual machines are no longer necessary.
- ended the war between users. Users agreed that machines partly controlled by GRD are going to be GRD-only machines. Previously, resources were degraded for everyone when a few users "overused" the system.
- gave system administrators the benefit of having resource allocation done automatically, as well as being able to monitor resources at both the site and individual job level.

These goals could be reached because of GRD's policy enforcement capabilities which are ground-breaking for a job management system: GRD maintains control over running jobs and controls their resource share by means of the Cray and SGI operating system ("dynamic scheduling"). Thus, computing shares can be defined in advance and fully enforced. The fair share mechanism cannot be compared to existing Cray features, because it is based on a far more complex paradigm.

GRD's policy capabilities include automated reactions to high-priority requirements, on-time termination for deadline jobs, weighting of individual machines with respect to different factors, a well-defined interface for manual override without policy changes and fair share.

5:00
How CUG Members Support Their Users

Barbara Horner-Miller, ARSC

In April a survey was sent to all CUG member sites concerning how the site performs the User Services task. The results of this survey will be presented.

11A The Power of the Engine: Users and the O/S

Chair: Terry Jones, GRUMMAN

9:00
The Age-Old Question of How to Balance Batch and Interactive

Barry Sharp, BCS

The UNICOS system is designed with both batch and interactive workload requirements in mind. However, in practice, the vanilla UNICOS system memory scheduler struggles to adequately balance these two distinct workloads, even with its wealth of scheduler tuning parameters. This paper presents a simple modification to the kernel-level memory scheduler that works harmoniously with the vanilla system to simplify the process.

9:30
MISER: User Level Job Scheduler

Nawaf Bitar, SGI/Cray

Miser is a dynamic batch process scheduler that provides deterministic batch scheduling and increased system throughput without static partitioning of a system. Static partitioning allows resources to be targeted at specific classes of processes, thus allowing deterministic run times, but can lead to waste. Miser

gives submitted batch jobs processing priority over the required resources needed to complete on schedule, but reduces waste by allowing idle resources to be used. The batch scheduler consists of a user process which schedules job resources, and kernel support for delivering resources to jobs on schedule. The combination allows flexibility while improving responsiveness and performance.

10:00

Serving a Demanding Client While Short on Resources

Manfred Stolle, ZIB

DMF on UNICOS, UNICOS/MK and IRIX is well suited to handle a local file system but it gets rather poor in a distributed environment with files larger than the DMF controlled file system. Using a small DMF controlled file system a "file system full" error is not supposed to occur due to poor migration performance or too large external files. A DMF controlled file system with a long history shows problems concerning the fast and safe restoration of data located on damaged tapes. Old style file attributes are disturbing in these situations. Solutions of these problems are shown.

11B Applications and Algorithms

Chair: Richard Shaginaw, BMSPRI

9:00

Application Roadmap

Greg Clifford, SGI/Cray

The SGI HPC computing environment will be advancing very rapidly in the next few years. This presentation will focus on how the SGI application environment will evolve and transition. The emphasis will be on solution areas (e.g., crash, external CFD) rather than on specific applications (e.g. MSC/NASTRAN, FLUENT).

9:30

Mathematical Methods for Mining in Massive Data Sets

Helene E. Kulsrud, CCR-P/IDA

With the advent of higher bandwidth and faster computers, distributed data sets in the petabyte range are being collected. The problem of obtaining information quickly from such data bases requires new and improved mathematical methods. Parallel computation and scaling issues are important areas of research. Techniques such as decision trees, vector-space methods, bayesian and neural nets have been applied. A short description of some successful methods and the problems to

which they have been applied will be presented. Methods which work effectively on PVP machines will be emphasized.

11C GigaRing Operations

Chair: Mike Brown, EPCC

9:00

GigaRing Configuring and Dumping

Michael Langer, SGI/Cray

A technical interchange regarding system configuration with multiple mainframes on the same GigaRing. The discussion would include SWS configuration options and administration issues. Single SWS support discussion will be included as well as a discussion on dumping GigaRing based systems.

9:30

Monitoring and Automatic Reboot of Cray GigaRing Systems

Birgit Naun, Thomas Plaga, KFA; Ralph Krotz, SGI/Cray

In the operatorless environment of a modern computer center the automatic reboot of production systems in trouble is an urgent requirement.

Abstracts

This paper describes the distributed monitoring mechanism on the Cray mainframes, the corresponding system workstations (SWS) and the computer center's problem management database server. If this monitoring detects a malfunction (could be any hardware or software problem) a set of utilities implemented on the SWS gathers as much information as possible for later problem analysis, reboots the failing component or the complete system if required, sends out Email messages and communicates status changes to the central database.

The automatic rebooting mechanism is designed to run on top of the Cray supplied basic system operation commands and supports all GigaRing based Cray systems (including T3E, T90 and J90se) and can be adapted easily to local information handling requirements.

13A Hauling Big Data

Chair: Hartmut Fichtel, DKRZ

2:00

Mass Storage at the NCSA: DMF and Convex UniTree

Michelle Butler, NCSA

The NCSA has been using UniTree since 1993 as the primary mass storage system for user's files and enterprise wide backups. The rapid evolution and growth of the NCSA production supercomputing environment, currently 512 Origin 2000 processors, a Power Challenge Array and an HP/Convex Exemplar, is putting increasingly critical requirements on the mass storage system. To meet this challenge, we are creating an environment which consists of a high performance mass storage system using DMF running on an Origin 2000 and an enterprise wide backup UniTree system on a recently installed HP/Convex Exemplar SPP-2000. This requires a transition from UniTree to DMF for our users and their data, approximately 15 TB. We will describe our experiences with this process using SGI/Cray tools for the migration as well as providing some comparison of the capabilities of the two systems.

2:30

Towards Petabytes of High Performance Storage at Los Alamos

Gary Lee, LANL

The High Performance Storage System (HPSS) is currently deployed in the open and secure networks at Los Alamos National Laboratory. Users of the Accelerated Strategic Computing Initiative (ASCI) Blue Mountain MPP system and the Advanced Computing Laboratory's MPP system, both from SGI/Cray, access HPSS for their storage. We discuss our current HPSS configurations, how our MPP users access HPSS, and the performance between HPSS and the MPP systems. We will also discuss our projected storage and storage performance requirements for the next five to seven years.

3:00

Storage and Data Management— Big Data Solutions

Ken Hibbard, SGI/Cray

This talk summarizes the storage road map for the Origin series of systems. In addition to describing existing and new hardware and software capabilities, we will highlight fiber channel solutions, file system support for new topologies, and data management solutions that are necessary for managing terabyte and petabyte stores. We will conclude with a summary of recent performance achievements.

Abstracts

13B Secure & Controlled Computing Engines

Chair: Bonnie Hall, LANL

2:00

Securing the User's Work Environment

Nick Cardo, SS-SSD

High performance computing at the Numerical Aerospace Simulation Facility at NASA Ames Research Center includes C90's, J90's and Origin 2000's. Not only is it necessary to protect these systems from outside attacks, but also to provide a safe working environment on the systems. With the right tools, security anomalies in the user's work environment can be detected and corrected. Validating proper ownership of files against user's permissions, will reduce the risk of inadvertent data compromise. The detection of extraneous directories and files hidden among user home directories is important for identifying potential compromises. The first runs of these utilities detected over 350,000 files with problems. With periodic scans, automated correction of problems takes only minutes. Tools for detecting these types of problems as well as their development techniques will be discussed with emphasis on consistency, portability and efficiency for both UNICOS and IRIX.

2:30

The State of Security for UNICOS & IRIX

Jay McCauley, SGI/Cray

I will review general security administration available for the UNICOS & IRIX operating systems. A review of the latest methods of breaking into systems and networks will be identified and discussed. A detailed analysis of password file security and password protection methods will be covered. Resources available both in printed form and "on the net" will be listed. Examples will be provided. Recommendations for increasing site security will be discussed in a short Question and Answer session after the presentation of the paper.

3:00

IRIX Accounting Limits and UDB Functionality

Jay McCauley, SGI/Cray

This talk will discuss how IRIX is being extended in the areas of accounting and limits to enable it to perform effectively in a data center environment.

13C User Services

Chair: Leslie Southern, OSC

2:00

Coordinating User Support Across a Partnership: NPACI User Services

Jay Boisseau, SDSC

The National Partnership for Advanced Computational Infrastructure (NPACI) includes 5 resource partners that provide computing resources (including CRAY PVPs and T3Es) for over 4000 users. NPACI is striving to provide the highest caliber of user services by combining and leveraging the staff expertise of the resource partners. Coordinating the efforts of different sites to develop synergy and not suffer from increased overhead demands careful planning, flexible procedures, and effective tools. We present the plans, procedures and tools developed by the NPACI Resources Working Group to offer consulting, training, documentation, and other user services to enable computational scientists to excel.

2:30

Cray Training Update

Bill Mannel, SGI/Cray

On May 2, 1998, Silicon Graphics Customer Education closed the Eagan, Minnesota train-

Abstracts

ing center. Future offerings for both IRIX and UNICOS-based curricula will be available in other worldwide training centers or at customer sites. This talk will discuss past and future changes in Cray training.

3:00

Pushing/Dragging Users Towards Better Utilization

Guy Robinson, ARSC

This talk will describe ARSC experiences in the operation of MPP Cray systems. In particular issues of user/center interaction will be considered with regard to:

- education in and use of the latest software tools and hardware features
- dealing with parallel experts and novices
- system configuration issues
- how to ensure the best use is made of the available resources by all users
- what to change in the system configuration and what/how to change users' practices
- ensuring users make full use of advanced computational methods, both in terms of system hardware/software developments and by exploiting suitable algorithms.

It is hoped that the talk will promote discussion of the need for centers to become more involved with users and how this can be achieved.

OPEN MEETINGS

4:00, 5:00, and 5:30

Please plan to attend the Open Meetings to discuss topics of interest to each Special Interest Group. The location of these meetings will be posted on the bulletin board.

Abstracts

15A IRIX: From Plans to Reality

Chair: Nick Cardo, SS-SSD

9:00

Getting It All Together

Daryl Grunau, LANL

At Los Alamos National Laboratory we have been operating two 256 processor clusters of Origin 2000's for over a year, recently expanding to 1280 processors. The primary programming model is SGI/MPI, with job sizes ranging up to 512 processors and routinely spanning machines. DFS is used for primary local storage, with HPSS targeted for archival storage. Free access to the nodes is restricted, and use of the resources is managed by Load Sharing Facility (Platform Computing). We have been active in pushing ahead the latest releases of the operating system, and currently have early experiences with IRIX 6.5 to report. The purpose of this paper will be to share the challenges, trials, and successes in getting it all to work together.

9:30

Integrating an Origin 2000 into a Cray Data Center

Chuck Keagle, BCS

This paper will present an overview of the Data Center configuration and discuss the application mix for an Origin 2000/Cray T916-256 Data Center. Major Topics will include:

- Hardware Configuration—Origin 2000, Cray T916-256, Connectivity
- Configuration Control—OS Versions, Patches, Local Code
- User Maintenance—UID Management, User Account Management
- Operations—Remote console operation using IRIS console
- Performance Tuning—Interactive vs. Batch limits, Performance Co-Pilot
- File System Structure—SCSI, Fiber RAID, DMF, mkfs/mount options
- Application Mix—Scalar codes, Vector codes, CPU availability

10:00

Getting the Best Mileage out of Your Origin System

Jim Harrell, SGI/Cray

How to turn various kernel knobs, allocate node-local data properly and use system profiling tools to get the best performance out of your Origin System.

15B Applications and Algorithms

Chair: Richard Shaginaw, BMSPRI

9:00

XVM–Extended Volume Management

Colin Ngam, SGI/Cray

Xtended Volume Management (XVM) is the next generation product to XLV. XVM provides extended functionalities in the IRIX and Cellular IRIX Operating System environment. XVM manages the following logical objects: Volumes, Subvolumes, Concats, Mirrors, Stripes and Physical Slices. XVM provides transparent distributed access to all XVM objects in both the Enterprise (cluster) and HPC (Cellular) environment. This paper provides a discussion on the differences between XVM and XLV, XVM and UNICOS/UNICOSmk Logical Device Drivers and the full set of XVM functionalities.

9:30

The Good, the Bad, and the Ugly Aspects of Installing New OS Releases

Barry Sharp, BCS

Many SGI/Cray sites do not have the luxury of having multiple like systems, one of which to serve the needs of validating a new OS release

Friday, June 19, 1998

prior to installing it as their production system. The difficulty of doing this on a single system without endangering the running production system is explored along with how Boeing has developed its processes over time to safeguard production while still being able to adequately QA the new OS release.

10:00

Origin Craylink Partitioning

Steve Whitney, SGI/Cray

Partitioning a machine involves dividing a single large machine into multiple machines, each with their own IP address, root disk, etc. with a very fast interconnect between them. A single partition can be brought down (S/W and H/W failures, controlled shutdown, etc.) without affecting the rest of the machines, thus providing higher availability for a partition than a single large machine.

15C Visualization

Chair: L. Eric Greenwade, INEEL

9:00

CUG Video Theater

Supercomputers generate data that, via scientific visualization, can be transformed into pictures that provide a way for scientists to understand the information more fully and quickly. Although the primary purpose of visualization is to facilitate the discovery process, scientists are increasingly coming to rely on it to also present their research conclusions. This session showcases the latest in scientific/engineering visualization. It promises to be both educational and entertaining.

10:00

Visualization of 3 Dimensional Material Science Applications

L. Eric Greenwade, INEEL

The INEEL has a number of programs investigating the properties of materials under demanding environmental conditions. These are related to a number of problems in environmental, military and industrial processing and reliability areas. This presentation will illustrate the results of some of these investigations and the new approaches used to convey the computational results to the researchers.

16 GENERAL SESSION

Parallel and Distributed Development and Simulation of Atmospheric Models

V. Mastrangelo and I. Mehilli, CNAM-Université; F. Schmidt, M. Weigele, J. Kaltenbach, A. Grohmann, and R. Kopetzky, University of Stuttgart

Parallel and distributed computing is a way to meet the increasing demand for engineering and computing power of scientific and technical simulation applications. To utilize these new paradigms the authors of this paper are developing a service based simulation environment which is operating on various computers at different European locations. In this paper we intend to discuss possibilities and difficulties of this approach by referring to the example of the dispersion of air carried particles. This example consists of different parts, where the air stream simulation (wind simulation) and the simulation of air carried particle transport (transport simulation) form the computationally expensive core of the application. This is run on a CRAY-T3E near Paris. To provide distributed calculation modules as a service various actions have to be taken. They include:

• Modularisation of the overall problem,
• Developing a strategy for distribution,

- Parallelisation and optimization of the computationally expensive
- Encapsulation of modules as independent processes,
- Developing a strategy for the integration of competing services into a simulation system,
- Performing a set of parallel simulations by providing the services of the system with consistent data and managing the distribution of the results.

Solutions to most of the problems will be proposed. A hierarchy of parallelisation steps will be discussed in connection with a corresponding strategy for distribution. Communication of parts with high granularity takes place on the basis of CORBA mechanisms. A special Service Agent Layer (SAL) provides methods to convert, modules or objects (with functionalities) into services available in a multi-user environment. Experiences with this architecture and the implementation of parts of it will be reported. Special emphasis is put on experiences including Cray computers into such a scenario.
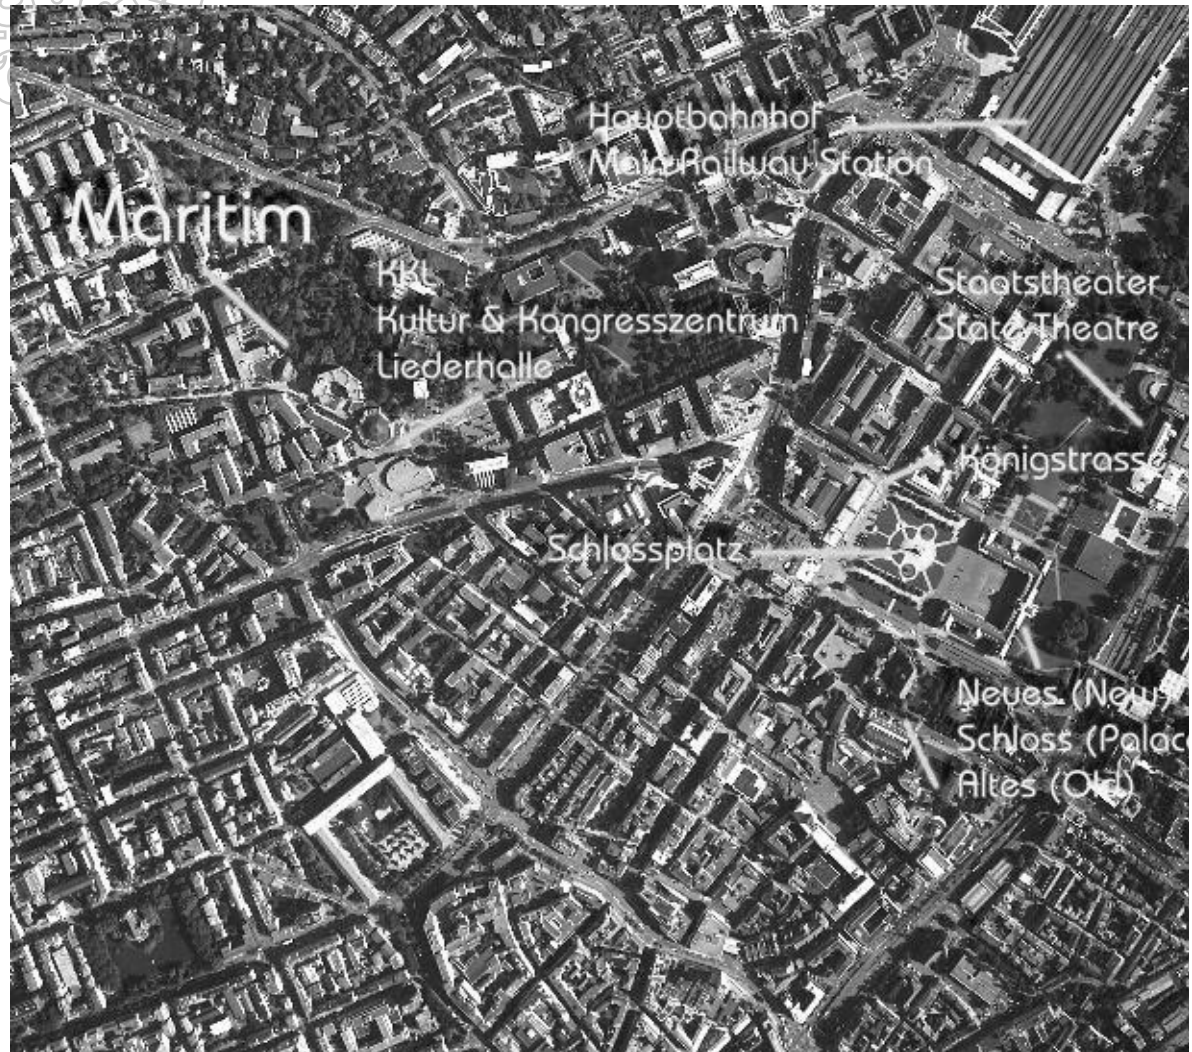
An aerial view of Stuttgart.

How to Contact Us

From June 13-19, 1998

Stuttgart CUG Conference
Grethe Knapp Christiansen
KKL
Berliner Platz 1-3
D-70174 Stuttgart
Germany
Phone: +49-(0) 711-2027.800 or 801
Fax: +49-(0) 711-2027.860
Email: stuttgart@cug.org

After the Conference

Grethe Knapp Christiansen
Institut für Computeranwendungen II
Universität Stuttgart
Pfaffenwaldring 27
D-70569 Stuttgart
Germany
Phone: +49-(0) 711-685.3759
Fax: +49-(0) 711-685.3758
Email: stuttgart@cug.org

Conference Information

Conference Site

KKL Stuttgart
Kultur- und Kongresszentrum Liederhalle
Berliner Platz 1-3
D-70174 Stuttgart
Phone: + 49-(0)711-2027.800 and
                           - 2027.801
Fax: + 49-(0)711-2027.860
Email: stuttgart@cug.org

In cooperation with: **Hotel Maritim**
Seidenstr. 34
D-70174 Stuttgart
Phone: + 49-(0)711-942.0
Fax: + 49-(0)711-942.1000

All sessions take place in the KKL. Lunches will be served in Saal Maritim on the ground floor in the adjacent Hotel Maritim.

Who May Attend?

CUG bylaws specify that employees of a CUG member (usually a facility or company using a Cray computer, identified by its CUG site code) and users of computing services provided by a CUG member, may attend a CUG meeting. Visitors are also eligible to attend a CUG meeting, provided they have formal approval from a member of the CUG Board of Directors. The members of the Board of Directors are listed later in this booklet.

Conference Registration Location/Hours

Location: First floor "Ebene 2" in KKL.

Office Hours:

Sun, June 14	3:00–5:00 p.m.
Mon–Thu, June 15–18	8:00 a.m.–6:00 p.m.
Fri, June 19	8:00 a.m.–1:00 p.m.

Badges and registration material may be obtained at these scheduled times. Note that ALL attendees must wear a conference badge when attending Stuttgart CUG Conference activities.

Call for Papers

There is a Call for Papers at the end of this brochure. To submit a paper for the Spring 1999 Cray User Group Conference in Minneapolis, Minnesota, USA, please complete and return the form as directed. All questions regarding the program should be directed to the Program Chair.

Yet another beautiful Platz!!

Local Arrangements

On Site Facilities

Throughout the conference, our staff at the CUG Conference Office will be glad to assist you with any special requirements.

Messages

During the conference (June 15-19) incoming telephone calls will be received at the CUG Conference Office Telphone: + 49-(0)711-2027.800 and -2027.801; Fax: + 49-(0)711-2027.860, Email: stuttgart@cug.org. Messages will be posted near the registration desk. Private telephone calls may be made in the downstairs entrance hall of the KKL using coins or telephone cards.

Faxes

A fax machine will be available at the CUG Conference Office for conference-related business.

Fax machines are available at the participating hotels for non-conference business. Please check with your hotel for prices for national and international fax transmissions.

Email and Personal Computers

The Email room is located in Donau (4.3.14). For last minute editing of papers and transparencies a Macintosh and a PC are available.

Photocopying

A copy machine will be available for making a limited number of copies. If you plan to distribute copies of your presentation materials, please bring sufficient copies with you. For nearby copy services (to be paid in cash only), please see the KKL Conference Office.

Dining Services

The conference fee includes refreshments during breaks and luncheons Tuesday through Thursday at the Hotel Maritim. Breaks will also be provided on Friday. Breakfast is included in the hotel price. Lunch will be served in the Hotel Maritim in Saal Maritim on the ground floor. On Monday, Newcomers are invited to a Newcomers' Lunch in the Hotel Maritim. Please sign up for this on the Conference Registration Form. On Monday evening, SGI/Cray Research will host a reception for all CUG attendees. On Wednesday evening, conference attendees are cordially invited by debis Systemhaus GmbH to a magnificent evening in the wonderful Mercedes-Benz Museum in Stuttgart.

Special accommodations and/or dietary requirements should be noted on the Conference Registration Form and/or on the Hotel Reservation Form under "Special requirements."

Hotel Information

Conference Hotel

Maritim Hotel Stuttgart

Seidenstr. 34
70174 Stuttgart
Phone: +49-711-942.0
Fax: +49-711-942.1000

Hotel Maritim is situated adjacent to the Conference center, KKL.

Other Hotel Options

A limited number of rooms have been reserved at these two hotels at a special weekend rate for the Saturday prior to the conference and the Friday after the conference in case you should wish to extend your visit in Stuttgart. Rooms are available on a first come first served basis. Special arrangements for families and students are available directly from Stuttgart-Marketing GmbH.

Best Western Hotel Ketterer
Marienstr. 3
70178 Stuttgart
Phone: +49-711-2039.0
Fax: +49-711-2039.600

Hotel Royal
Sophienstr. 35
70178 Stuttgart
Phone: +49-711-62505.0
Fax: +49-711-628809

Travel Information

Passports

A valid passport is required for travel to Europe and between countries once here. Please check with your travel agent or your local German consulate to determine if you need a visa and for information on how to obtain one. For E.C. residents a national ID is sufficient.

Transportation

Stuttgart is well equipped with official transport services. In the City, these are underground city railways, trams, and buses. Your conference registration packet includes a three day ticket for the Stuttgarter Straßenbahn and buses–VVS. This ticket is valid for one person for three consecutive days, including the first day of use. You may buy additional tickets in your hotel. Additonal information about the VVS is included in your conference registration packet.

We suggest taking a taxi from your hotel to the railway station (Hauptbahnhof) and then take the suburban railways (S2 and S3) to the Stuttgart airport. Trains leave approximately every 15 minutes and the trip takes approximately 27 minutes. You buy your ticket at an automatic ticket machine in the railway station for your trip to the airport. It costs DM 4,60. The taxi fare from downtown Stuttgart to the airport costs approximately DM 50,00.
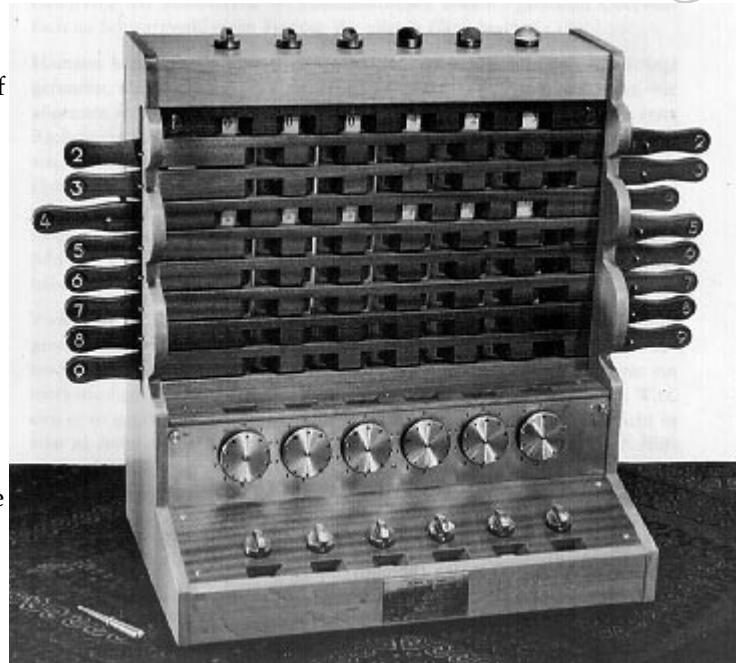
Currency

German currency is the Deutsche Mark (DM). Currency can be changed and traveler's checks can be cashed in banks which are generally open Monday to Friday from 9:00 a.m. to 12:00 p.m. and 2:00 p.m. to 4:00 p.m. and on Thursday to 5:30 p.m. Exchange offices are open every day at all major railway stations, at airports, and at border crossing points.

Voltage

The standard voltage in Germany is 230 volts, 50 cycles. Bring an adapter plug for use in electrical outlets.

Smoking

Smoking is not allowed in conference locations (except outside).

Special Assistance

Session rooms for the Stuttgart CUG Conference are located on several floors in the KKL and are accessible by stairs or elevators. If you need special assistance, please come to the Conference Office where our staff will be pleased to assist you.

Climate

Stuttgart is situated in the south of Germany and the climate is semicontinental, making it hard to tell what the weather will be like. Although the month of June is usually a nice summer month with sunny weather and comfortable temperature 18 to 24° C (65°–75°F), it may also be rainy and "chilly." To be safe, we suggest that you bring an umbrella.

Social Events

Monday, June 15

Newcomers' Luncheon

First-time CUG participants are invited to the Newcomers' Luncheon. At this event, newcomers can meet the CUG Board of Directors, the Special Interest Group Chairs, and executives from SGI/ Cray Research. This will also be a chance to hear a very brief presentation of "What Is CUG?". Please be sure to sign up for this luncheon at the Conference Office.

SGI/Cray Research Reception

All conference participants and their guests are invited by SGI/Cray Research to attend a reception on Monday, June 15, from 7:00–10:00 p.m in the Old Riding Room of the Maritim Hotel.

Wednesday, June 17

CUG Night Out

The traditional CUG night out will be held Wednesday, June 17, 1998. Our conference partner, debis Systemhaus GmbH, has invited all participants and their guests to the Mercedes-Benz Museum in Untertürkheim, the traditional place of the Mercedes-Benz factory. Here we shall dine and stroll around the Mercedes Benz Oldtimers in admiration of the "good old days." We are very happy for this invitation which will be a highlight of the Stuttgart CUG Conference. Trans-

portation will be provided to and from the museum by debis.

Additional guests must be noted on your Conference Registration Form for this event.

Site Tours

Tours of the Cray computer facilities at RUS can be arranged on short notice during the Conference. Please come to the Conference Office to make arrangements.

Accompanying Persons Program

Within the registration area at KKL, a counter with tourist information will be arranged where you may collect information about Stuttgart and Baden-Württemberg.

Why not stroll in the city just to get acquainted and see Königstrasse, the Palace Gardens, the State's Theaters, Schlossplatz, Marketplace, shop at Calwer Passage or just relax and try some Maultaschen and a glass of wine.

- The **State Gallery** in the center of Stuttgart with its impressive Picasso collection is closed on Monday so you may want to plan this for some other day during the week.

- The **Wilhelma**, Germany's only zoological and botanical garden, is situated near the river Neckar also not far from the city.

- The city of **Ludwigsburg** with its castle known as the **"Swabian Versailles"** and the famous **"Blooming Baroque,"** the castle's park (20 min. by public transportation).

- The city of **Esslingen**, with its well preserved medieval city center and its old town hall with an impressive renaissance facade and delightful carillon (20 min. by public transportation).

The **State Gallery**

Contacts

CUG Board of Directors

President

Gary Jensen
National Center for Supercomputing
  Applications
152 Computing Applications Building
605 East Springfield Avenue
Champaign, IL 61820 USA
Phone: (1-217) 333-1189
Fax: (1-217) 244-1987
guido@ncsa.uiuc.edu

Vice-President

Sam Milosevich
Eli Lilly and Company
Computational Science Apps, MC499
Lilly Corporate Center
Indianapolis, IN 46285 USA
Phone: (1-317) 276-9118
Fax: (1-317) 277-3565
sam@lilly.com

Secretary

Gunter Georgi
Northrop Grumman
Data Systems and Services Division
73 Irma Avenue
Port Washington, NY 11050 USA
Phone: (1-516) 883-2336
Fax: (1-516) 883-4823
georgi@cug.org

Treasurer

Barbara Horner-Miller
Arctic Region Supercomputing Center
910 Yukon Drive, Suite 108, P.O. Box 756020
Fairbanks, AK 99775-6020   USA
Phone: (1-907) 474-5409
Fax: (1-907) 474-5494
horner@arsc.edu

Director—Asia/Pacific

Shigeki Miyaji
Chiba University
Dept. of Physics
1-33, Yayoi-cho, Inage-ku
263 Chiba, Japan
Phone: (81-43) 290-3719
Fax: (81-43) 290-3720
miyaji@sirius.c.chiba-u.ac.jp

Director—Europe

Walter Wehinger
Rechenzentrum Universitaet Stuttgart
Allmandring 30
D-70550 Stuttgart
Germany
Phone: (49-711) 685-2513
Fax: (49-711) 682-357
wehinger@rus.uni-stuttgart.de

Director—Americas

Barry Sharp
Boeing Shared Services Group
Engineering Operating Systems
MC 7J-04
P.O. Box 3707
Seattle WA 98124-3707 USA
Phone   (1-425) 865-6411
Fax: (1-425) 865-2007
bxs@sdc.cs.boeing.com

Past President

Claude Lecoeuvre
Commissariat a l'Energie Atomique/
  Ile de France
DRIF/DP2I/EC
Centre de Bruyeres-le-Chatel, BP n 12
F-91680 Bruyeres-le-Chatel, France
Phone: (33-1) 45 95 61 85
Fax: (33-1) 69 26 70 69
clecoeuvre@cea.fr
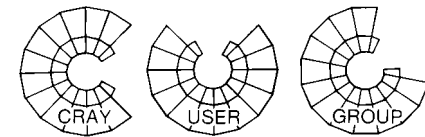
Special Interest Groups

Applications and Algorithms

Chair: Richard Shaginaw

Bristol-Myers Squibb Pharmaceutical
   Research Institute
P.O. Box 4000, Mail Stop J22-02
Princeton, NJ 08543-4000 USA
Phone: (1-609) 252-5184
Fax: (1-609) 252-6163
shaginaw@bms.com

Deputy: Larry C. Eversole

Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive, MS 126-147
Pasadena, CA 91109-8099 USA
Phone: (1-818) 354-2786
Fax: (1-818) 393-1187
larry.c.eversole@jpl.nasa.gov

Graphics

Chair: Eric Greenwade

INEEL-Lockheed Martin Idaho Technologies Co.
Lockheed Martin Idaho Tech.
P.O. Box 1625, MS-3605
Idaho Falls, ID 83415 USA
Phone: (1-208) 526-1276
Fax: (1-208) 526-4017
leg@inel.gov

SGI/Cray Contact: Cal Kirchhof

Visualization Program Manager
655F Lone Oak Drive
Eagan, MN 55121 USA
Phone: (1-612) 683-5245
Fax: (1-612) 683-5599
cal@cray.com

IRIX

Chair: Nicholas Cardo

Sterling Software, Inc.
NASA Ames Reseach Center
MS 258-6
Moffett Field, CA 94035-1000 USA
Phone: (1-415) 604-4754
Fax: (1-415) 964-1760
cardo@nas.nasa.gov

J90

Chair: Jean Shuler

Lawrence Livermore National Laboratory
Livermore Computing
P.O. Box 808, L-67
Livermore, CA 94552 USA
Phone: (1-510) 423-1909
Fax: (1-510) 422-0592
jshuler@llnl.gov

SGI/Cray Contact: Michelle Webster

SGI/Cray Research
Customer Service
655F Lone Oak Drive
Eagan, MN 55121 USA

Phone: (1-612) 683-5301
Fax: (1-612) 683-5307
maw@cray.com

Mass Storage Systems

Chair: Robert J. Silvia

North Carolina Supercomputing Center
3021 Cornwallis Road
P.O. Box 12889
Research Triangle Park, NC 27709 USA
Phone: (1-919) 248-1132
Fax: (1-919) 248-1101
rjs@ncsc.org

Deputy: Helene E. Kulsrud

Institute for Defense Analyses
CCR-P/IDA
Thanet Road
Princeton, NJ 08540 USA
Phone: (1-609) 279-6243
Fax: (1-609) 924-3061
laney@ccr-p.ida.org

European Deputy: Hartmut Fichtel

Deutsches Klimarechenzentrum GmbH
Systems Software
Bundesstrasse 55
D-20146 Hamburg 13, Germany
Phone: (49-40) 41-173-220
Fax: (49-40) 41-173-270
fichtel@dkrz.de

Contacts

Networking

Chair: Hans Mandt

Boeing Advanced Systems Laboratory
P.O. Box 24346, MS 7L-48
Seattle, WA 98124-2207 USA
Phone: (1-206) 865-3252
Fax: (1-206) 865-2965
hans@skipper.rt.cs.boeing.com

European Deputy: Dieter Raith

Rechenzentrum der Universitaet Stuttgart
Systems LAN
Allmandring 30
D-70550 Stuttgart, Germany
Phone: (49-711) 685-4516
Fax: (49-711) 682-357
raith@rus.uni-stuttgart.de

SGI/Cray Contact: Michael Langer

Manager, I/O Technical Computing
SGI/Cray Research
655F Lone Oak Drive
Eagan, MN 55121 USA
Phone: (1-612) 683-5801
mlanger@cray.com

Operating Systems

Chair: Terry Jones

Northrop Grumman Data Systems and Services
Bldg 1001, Room 101
Stennis Space Center, MS 39522 USA
Phone: (1-601) 688-5289

Fax: (1-601) 689-0400
jonestl@navo.hpc.mil

Deputy: Virginia Sotler

United States Army Corps of Engineers
Waterways Experiment Station
3909 Halls Ferry Road, MS CEWES-IM-H
Vicksburg, MS 39180-6199 USA
Phone: (1-601) 634-4418
Fax: (1-601) 634-2301
sotler@wes.army.mil

Operations

Chair: Daniel Drobnis

San Diego Supercomputing Center
P.O. Box 85608
SDSC-103
San Diego, CA 92186-5608 USA
Phone: (1-619) 534-5000
Fax: (1-619) 534-5152
drobnisd@sdsc.edu

Deputy: Virginia Bedford

Arctic Region Supercomputing Center
University of Alaska
P.O. Box 756020
Fairbanks, AK 99775-6020 USA
Phone: (1-907) 474-5426
Fax: (1-907) 474-5494
bedford@arsc.edu

European Deputy: Michael Brown

Edinburgh Parallel Computing Centre
University of Edinburgh, JCMB
King's Buildings, Mayfield Road

Edinburgh, Scotland EH9 3J2 UK
Phone: (44-131) 650 5031
Fax: (44-131) 650 6555
m.w.brown@ed.ac.uk

Performance and Evaluation

Chair: Jeff Kuehn

National Center for Atmospheric Research
Scientific Computing Divisioin
P.O. Box 3000
Boulder, CO 80307-3000 USA
Phone: (1-303) 497-1311
Fax: (1-303) 497-1804
kuehn@ucar.edu

Deputy (Americas): Bruce Kelly

Lawrence Livermore National Laboratory
L-73
PO Box 808
Livermore, CA 94551 USA
Phone: (1-510) 423-0640
kellybe@llnl.gov

Deputy (Europe): Michael Resch

High Performance Computing Center
  Stuttgart
Allmandring 30
D-70550 Stuttgart Germany
Phone: (49-711) 6855834
Fax: (49-711) 6788363
resch@rus.uni-stuttgart.de

SGI/Cray Contact: Jeff Brooks

SGI/Cray Research
Software Development
655F Lone Oak Drive
Eagan, MN 55121 USA
Phone: (1-612) 683-3410
jpb@cray.com

Security

Chair: Bonnie Hall

Los Alamos National Laboratory
Group CIC-3, MS B265
Los Alamos, NM 87545 USA
Phone: (1-505) 667-0215
Fax: (1-505) 665-5220
bhall@lanl.gov

SGI/Cray Contact: Jim Grindle

SGI/Cray Research
Software Development
655F Lone Oak Drive
Eagan, MN 55121 USA
Phone: (1-612) 683-5599
Fax: (1-612) 683-5599
jsg@cray.com

Software Tools

Chair: Hans-Hermann Frese

Konrad Zuse-Zentrum für
   Informationstechnik Berlin
Supercomputing Department
Takustrasse 7
D-14195 Berlin-Dahlem Germany
Phone: (49-30) 84185-145
Fax: (49-30) 84185-311
frese@zib.de

Deputy: Victor Hazlewood

San Diego Supercomputer Center
P.O. Box 85608
SDSC-0505
San Diego, CA 92186-9784 USA
Phone: (1-619) 534-5115
Fax: (1-619) 534-5077
victor@sdsc.edu

SGI/Cray Contact: Bill Harrod

SGI/Cray Research
Software Development
655F Lone Oak Drive
Eagan, MN 55121 USA
Phone: (1-612) 683-5249
Fax: (1-612) 683-5276
harrod@cray.com

User Services

Chair: Barbara Horner-Miller

Arctic Region Supercomputing Center
910 Yukon Drive, Suite 108
P.O. Box 756020
Fairbanks, AK 99775-6020 USA
Phone: (1-907) 474-5409
Fax: (1-907) 474-5494
horner@arsc.edu

Deputy: Leslie Southern

Ohio Supercomputer Center
User Services
1224 Kinnear Road
Columbus, OH 43212-1163 USA
Phone: (1-614) 292-9248
Fax: (1-614) 292-7168
leslie@osc.edu

SGI/Cray Contact: Mike Sand

SGI/Cray Research
Application Services
890 Industrial Blvd.
Chippewa Falls, WI 54729 USA
Phone: (1-715) 726-5109
Fax: (1-715) 726-4343
mas@cray.com

**SILICON GRAPHICS
DEUTSCHLAND**

Program Steering Committee

Chair: Sam Milosevich

Eli Lilly and Company
Computational Science Apps, MC499
Lilly Corporate Center
Indianapolis, IN 46285 USA
Phone: (1-317) 276-9118
Fax: (1-317) 277-3565
sam@lilly.com

Mary Kay Bunde

SGI/Cray Research
Software Development
655F Lone Oak drive
Eagan, MN 55121 USA
Phone: (1-612) 683-5655
Fax: (1-612) 683-5098
mkay@cray.com

Eric Greenwade

INEEL-Lockheed Martin Idaho Technologies Co.
Lockheed Martin Idaho Tech.
P.O. Box 1625, MS-3605
Idaho Falls, ID 83415
USA
Phone: (1-208) 526-1276
Fax: (1-208) 526-4017
leg@inel.gov

Bonnie Hall

Los Alamos National Laboratory
Group CIC-3, MS B265
Los Alamos, NM 87545 USA
Phone: (1-505) 667-0215
Fax: (1-505) 665-5220
bhall@lanl.gov

Barbara Horner-Miller

Arctic Region Supercomputing Center
910 Yukon Drive, Suite 108
P.O. Box 756020
Fairbanks, AK 99775-6020 USA
Phone: (1-907) 474-5409
Fax: (1-907) 474-5494
horner@arsc.edu

Helene Kulsrud

Institute for Defense Analyses
CCR-P/IDA
Thanet Road
Princeton, NJ 08540
USA
Phone: (1-609) 279-6243
Fax: (1-609) 924-3061
laney@ccr-p.ida.org

Conference Chair

Walter Wehinger

Rechenzentrum der Universitaet Stuttgart
Allmandring 30
D-70550 Stuttgart Germany
Phone: (49-711) 685-2513
Fax: (49-711) 682-357
wehinger@rus.uni-stuttgart.de

CUG Office

Bob Winget

2911 Knoll Road
Shepherdstown, WV 25443
USA
Phone: (1-304) 263-1756
Fax: (1-304) 263-4841
cug@cug.org

Notes on the Conference Room Names

We would like to make you aware of the rich historic background of the State of Baden-Württemberg that is captured by our conference room names.

Hohenzollern is a crown atop a forested hill, its many turrets reminiscent of knights in shining armour, fair princesses, and kings upon their thrones. If you like the fairy tale image of a castle, Hohenzollern should be on your itinerary.

Monrepos An enormous Baroque palace built at the beginning of the 18th century by the Duke of Württemberg, dreaming of another Versailles, was the start of Ludwigsburg—a town whose artificial character is emphasized by its orthogonal grid of dead end streets. To the Northwest of Ludwigsburg, nature and culture meet in perfect harmony at Schloß Monrepos. Built in 1767 on a site bordering a small lake, this gracious Rococo palace encircles a beautiful rotunda.

Solitude is located on a wooded elevation west of the city. Despite empty state coffers, Duke Karl Eugen gave the green light to build Schloß Soli-tude. He wanted brilliance, not quality. Therefore, the estate had to be restored joint by joint from 1975—1983 so that it could continue to display its breathtaking effect.

Donau or Fluvius Danuvius by its Roman name, German Donau, Slovak Dunaj, Serbo-Croatian and Bulgarian Dunav, Russian Dunay, this is Europe's second longest and most powerful river. Rising in the Black Forest mountains of Germany, it flows for approximately 2850 km through the southeastern portion of the continent to the Black Sea.

The **Neckar** river unites the landscapes of Württemberg and Baden on its way from the Black Forest to Mannheim, where it joins the Rhine river. Downstream from Bad Wimpfen, the Neckar, cutting once more through the sandstone massif of the Odenwald, runs through an area of wooded hills, many of them crowned by castles ... performing all of these efforts only to meet a true love—old Heidelberg.

Rhein The Rhine river, 1320 km long, flows through four different countries and as such has often been a source of controversy between neighbors. At the same time it has always been a unique highway for the exchange of commercial, intellectual, artistic and religious ideas.
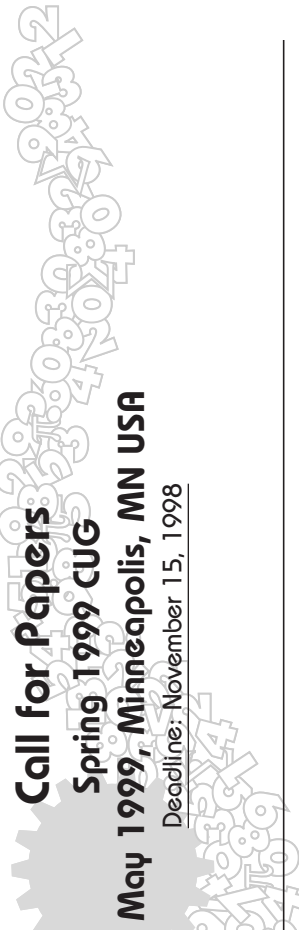
Schiller Marbach am Neckar, this small town cultivates the memory of the poet Friedrich Schiller (1759-1805), its most illustrious son. The huge Schiller-Nationalmuseum rising like a castle above the valley, displays portraits, casts, engravings and manuscripts. A visit to the writer's birthplace completes the pilgrimage.

Cray User Group Conference Stuttgart
Conference Building

Conference Room	4.4.20	Solitude			Ebene 4
	4.4.22-24	Hohenzollern			
	4.4.25-27	Monrepos			Level 4
Conference Room	4.3.10	Neckar	Email		Ebene 3
	4.3.11-13	Rhein			
	4.3.14-16	Donau	Demo/ Slide Check		Level 3
CUG Conference Office		Schiller-Foyer			Ebene 2
Main Entrance					Level 2
Coffee Breaks		Schiller-Foyer			Ebene 1
General Sessions		Schiller-Saal			
Access Hotel Maritim					Level 1

Elevator   Stairs

Call for Papers:
Spring 1999 CUG
May 1999, Minneapolis, MN USA

Deadline: November 15, 1998

Title of Paper

Abstract (two or three sentences)

Name

Organization

CUG Site Code(Mandatory)

Address

City   State/Province   Mail/Zip Code   Country

Daytime Phone

Fax

(Country code-city/area code) number   (Country code-city/area code) number

Electronic Mail Address   Person submitting paper's Email address

Audio/Visual Requirements
- ☐ Video 1/2"
- ☐ 35-mm slide projector
- ☐ Video Projector to attach to your laptop
- ☐ Video 3/4"
- ☐ Overhead projector
- ☐ Other?

Suggested Session
- ☐ General Session
- ☐ Applications and Algorithms
- ☐ Graphics
- ☐ IRIX
- ☐ J90
- ☐ Joint Session (list suggested sessions)
- ☐ Other
- ☐ Mass Storage Systems
- ☐ Networking
- ☐ Operating Systems
- ☐ Operations
- ☐ Performance and Evaluation
- ☐ Security
- ☐ Software Tools
- ☐ User Services
- ☐ Tutorial

Please list any co-authors' names and Email addresses

Please check the CUG home page for guidelines for preparing and submitting your paper: **www.cug.org**
You will be contacted by your Session Organizer to notify you of acceptance or rejection of your paper. If you do not hear from your
Session Organizer by January 15, 1999, then please contact Sam Milosevich, Program Steering Committee Chair, at the address below.

Return Address:

Sam Milosevich
Eli Lilly and Company
Computational Science Apps, MC499
Lilly Corporate Center, Drop 0715
Indianapolis, IN 46285-1513   USA
Phone: (1-317) 276-9118   Fax: (1-317) 277-3565
sam@lilly.com

CRAY USER GROUP

INCORPORATED