Proceedings of SRS 94-95[1]

Oregon Graduate Institute of Science & Technology
Department of Computer Science and Engineering
Student Research Symposium

January 27, 1995

[1]OGI Technical Report CSE-95-005

From the SRS 94-95 Committee ...

The 1994-95 Student Research Symposium (SRS 94-95) of the Oregon Graduate Institute's Department of Computer Science and Engineering was held on January 27th, 1995, at the OGI campus. The symposium was held to showcase the diverse research being pursued within the department by its students. We intended the day to provide opportunities for learning and exchange amongst both new and advanced students as well as our faculty and outside attendees, and it was clearly a success.

We were pleased with the student response to this Symposium, which resulted in a full day of eleven presentations. We thank all the students who contributed their work, and we congratulate the recipients of the awards:

Best Presentation:	Bennet Vance presented *New Techniques for Optimizing Joins of Many Relations*
Runner up:	Brian Mak presented *Two Perceptual Experiments for Continuous Speech Understanding*

Bennet was the recipient of the Touring Award, the new SRS traveling trophy.

The committee would like to thank the presenters, judges, audience, and everyone who helped.

Your comments and criticisms on the Symposium are welcome; feel free to speak to one of us, or email us at `ogi-srs@cse.ogi.edu`. Thanks again to all who participated in one way or another!

All information on the Symposium is available on the World Wide Web at the URL `http://www.cse.ogi.edu/SRS/94-95/`. This technical report is also available via anonymous ftp from `cse.ogi.edu` in `pub/tech-reports/1995/95-005.ps.gz`.

Symposium Committee
E. David Burke, *Symposium Chair*
Robert M. Prouty

SRS 94-95 Schedule
January 27th, 1995

10:30 Session 1
Opening Ceremonies
Andrew Burke, *Parallel Radiosity*
Tito Autrey, *Run Time Code Generation: What Is It?*
Jon Inouye, *MBONE and NERO: Providing Desktop Audio/Visual Tools to OGI*

noon *Lunch break*

1:00 Session 2
Priyadarshan Kolte, *Nascent: a Compiler for a Multithreaded, Cache-coherent, Multiprocessor Architecture*
Bennet Vance, *New Techniques for Optimizing Joins of Many Relations*
L. Don Colton, *False Alarm Rejection in Keyword Spotting Systems*

2:05 *Short break*

2:15 Session 3
Brian Mak, *Two Perceptual Experiments for Continuous Speech Understanding*
Richard Staehli, *Quality of Service Specification for Multimedia Presentations*
Daniel Burnett, *Speaker Adaptation through Modified Backpropagation Training*

3:20 *Short break*

3:30 Session 4
Lisa Walton, *Design Automation: Making Formal Methods Relevant*
Dan L. Clark, *Parallel Job Scheduling on Dynamic Heterogeneous Networks*
A few words about a local ACM chapter

4:30 Closing Remarks
Awards Banquet

Parallel Radiosity

Presented by Andrew Burke
aburke@cse.ogi.edu

Radiosity is a computer graphics rendering method specifically developed for rendering diffuse surfaces, and can be contrasted with ray tracing, which is more suited for rendering smooth shiny surfaces. Radiosity was introduced in the mid 1980's, and is well known for being an exceptionally slow and expensive method for generating images. This talk will describe a parallel implementation of radiosity. This talk will also include a short description of the toolset implemented for simplifying the development of parallel and distributed applications.

Run Time Code Generation: What is it?

Presented by Tito Autrey
tito@cse.ogi.edu

Run time code generation (RTCG) is a very hot topic in the OS and compiler fields. RTCG is where new code is generated and integrated into an executing program. An application can derive several benefits from RTCG: at run time a very general module can be highly specialized for improved performance, and new functionality can be added without halting an executing program. This talk will cover where to use RTCG, how to use it, and some of the issues to address for automatic RTCG.

MBONE and NERO: Providing Desktop Audio/Visual Tools to OGI

Presented by Jon Inouye
jinouye@cse.ogi.edu

The advent of high speed networking has allowed the Internet to provide additional services beyond that of the traditional email and file transfer. One such service, audio/visual conferencing, will facilitate the transfer of information such as classes and colloquia.

This talk describes the Multicast Backbone (MBONE) and the Network for Education and Research in Oregon (NERO). MBONE multicasts provide one-to-many and many-to-many packet delivery for applications such audio and video conferencing tools. NERO provides the high-speed network infrastructure required by some of these applications. Half of the talk will focus on how MBONE and NERO will influence OGI. The remainder of the talk will explore the technical aspects of MBONE and NERO.

Nascent: a Compiler for a Multithreaded, Cache-coherent, Multiprocessor Architecture

Presented by Priyadarshan Kolte
pkolte@cse.ogi.edu

This talk describes work in the Sparse compiler group to compile parallel Fortran programs for the *T-NG architecture. We shall first briefly cover the interesting characteristics of this architecture and then present an overview of the proposed compiler projects. These projects include compiler-based multithreading, compiler-directed cache coherence, automatic data and computation placement, interprocedural analyses, and late-binding optimizations.

New Techniques for Optimizing Joins of Many Relations

Presented by Bennet Vance
bennet@cse.ogi.edu

Join optimization—the process of finding cheap evaluation plans for multiway joins—is essential to efficient query processing. The classical solution, exhaustive search by way of dynamic programming, offers several advantages over competing techniques, not the least of which is that it guarantees optimality according to a given cost model. On the down side, exhaustive search has exponential complexity, and is reputedly intractable for joins of more than about 8 relations. However, this reputation is undeserved. In this talk we describe an approach to exhaustive search that can easily accommodate 15 relations, and for some classes of joins, 20 relations or more. Our approach is sufficiently simple and flexible that we believe it can be put to practical use in commercial database systems.

False Alarm Rejection in Keyword Spotting Systems

Presented by Don Colton
ldcolton@cse.ogi.edu

Given a putative hit (hypothesized utterance) found by a Keyword Spotting (KWS) System, our research is to reliably reject false alarms, thus improving overall recognition accuracy. Our technique is to examine the high-scoring phoneme in each time frame, and compare it and its score to the match used by the putative hit.

Our typical KWS is a first-pass keyword spotter that scans incoming utterances for specific phrases from a list. The spotter returns the best-scoring match when in fact the talker may have made an out-of-vocabulary utterance. In this case the best match is actually a false alarm.

Speech recognition software can be trained using many samples of each word to be detected. Such vocabulary-dependent training improves recognition confidence. But sufficient data is not available for large-vocabulary training. The goal of this research is to improve recognition confidence without vocabulary-dependent training.

This work is joint with Mark Fanty.

Two Perceptual Experiments for Continuous Speech Understanding

Presented by Brian Mak
mak@cse.ogi.edu

In this paper, we present two perceptual experiments designed to determine which parts of fluent speech are more important for human to understand the overall speech:

- vowels or consonants?
- transition segments or stable segments?

Sixty read sentences were selected from the TIMIT database, and they were manipulated for our experiment purposes. Seventy subjects participated in the evaluation, 35 for each experiment. Results show that vowels are more important than consonants in human understanding of continuous speech, while transition segments and stable segments are equally important.

Quality of Service Specification for Multimedia Presentations

Presented by Richard Staehli
staehli@cse.ogi.edu

Digital multimedia systems are rapidly becoming ubiquitous as support for audio and video is now available on nearly all computing platforms. Multimedia computing promises to augment or replace most of the traditional broadcast and print media with more interactive and personalized information services. Unfortunately, these existing systems still suffer from immature technology that offers limited inter-operability and connectivity. Two significant research problems are:

- real-time scheduling of shared and distributed resources
- information loss-management under overload conditions

In this talk I show that poor understanding of multimedia presentation requirements has resulted in ad hoc solutions to these problems. A new framework is proposed that allows formal specifications of presentation quality of service (QOS) requirements. These QOS specifications can be used to select optimal resource scheduling and overload handling policies.

Speaker Adaptation through Modified Backpropagation Training

Presented by Daniel Burnett
burnett@cse.ogi.edu

Psychoacoustic experiments have indicated that the human ear has better frequency sensitivity at lower frequencies than at higher frequencies. To account for this, some spoken language system designers warp frequencies present in the original speech signal to that of the Bark scale, which tends to compress the higher frequencies and expand the lower ones. However, because different speakers have different vocal tract characteristics, the same speech sound produced by different speakers may result in different initial frequency distributions.

One possible way of dealing with this difference is to adjust the Bark-scale warping to be different for each speaker. Other researchers have indicated that changing the rate at which the speech was digitally sampled will effectively change this warping. Unfortunately, this requires one to run a speech recognizer multiple times on the same utterance, while varying only the sampling rate, to determine which sampling rate is the best for the given speaker (for that utterance). We feel that a more efficient method of determining this sampling rate is to use the output error of our neural network classifier during backpropagation training. We will then be able to use the sampling rate determined from looking at one (or a small number of) utterance(s) from a speaker to adapt the inputs to our classifier to give improved performance for any following utterances from the same speaker.

Design Automation: Making Formal Methods Relevant

Presented by Lisa Walton
walton@cse.ogi.edu

Formal methods are easily dismissed as heavy weight and irrelevant. People legitimately ask to see what impact formal methods have had on software development practices outside of those areas where governments or regulatory agencies have mandated their use. This paper argues that design automation is an opportunity for the insertion of formal methods into software development practice. In particular, we advocate the development of small, domain-specific specification languages, and the use of these languages as front-ends to program generation systems.

Software Design for Reliability and Reuse (SDRR) is a method developed at the Pacific Software Research Center to support the development of generators for domain-specific languages. Our method, which utilizes a robust suite of reusable transformation tools, has been used to implement a software component generator for a message translation and validation problem domain identified by the Air Force. Independent contractors are conducting an experiment comparing our generation based technology to an existing solution that uses a program-templates based reuse technology.

Parallel Job Scheduling on Dynamic Heterogeneous Networks

Presented by Dan L. Clark
dclark@cse.ogi.edu

In using a shared network of workstations for parallel processing, it is not only important to consider heterogeneity and differences in processing power between the workstations but also the dynamics of the system as a whole. In such a computing environment where the use of resources vary as other applications consume and release resources, intelligent scheduling of the parallel jobs onto the available resources is essential to maximize resource utilization. Despite this realization, however, there are few systems available that provide an infrastructure for the easy development and testing of these intelligent schedulers. In this paper, an infrastructure is presented together with a scheduler that is capable of both gang scheduling and dynamic task reallocation of PVM applications.