

THE UNIVERSITY OF MICHIGAN
COMPUTING RESEARCH LABORATORY

Bibliography of Computing Research
Laboratory Reports, 1982-1983

CRI-TR-1-84

edited by
Virginia L. Folsom
and
Elizabeth I. Olsen

JANUARY 1984

Room 1079, East Engineering Building
Ann Arbor, Michigan 48109
USA
Tel: (313) 763-8000

enign
UMR1011

Bibliography
of
Computing Research Laboratory
Reports, 1982-1983

edited by
Virginia L. Folsom and Elizabeth I. Olsen
Publications Coordinators

January 1, 1984

Abstract:

This report lists, in chronological order, all reports published by the Computing Research Laboratory of the University of Michigan since its inception in 1982. Each report is identified by the Laboratory number, author(s), title, number of pages and date.

In the Bibliography which follows, there is a listing of abstracts of publications produced by the Computing Research Laboratory at The University of Michigan as of the date of this writing. These publications are available from

Computing Research Laboratory

The University of Michigan

1079 East Engineering

Ann Arbor, Michigan 48109

(313) 763-8000.

CRL-TR-1-82.

C.M. Krishna and K.G. Shin

"Performance Measures for Multiprocessor Controllers," 25 pages, October 1982.

abstract -

In this report, we consider some new performance measures to characterize fault-tolerant multiprocessors used in the control of critical processes. Our performance indices are based on controller response time. By relating this to the needs of the application, we have been able to derive indices that faithfully reflect the performance of the multiprocessor *in the context of the application*, that permit the design, evaluation and objective comparison of rival computer systems, and that can either be definitively estimated or objectively measured.

Using the example of a controller in an idealized satellite application, the computation, and some uses of, the performance measures are illustrated.

CRL-TR-2-82.

Kang G. Shin and Yann-Hang Lee

"Analysis of the Impact of Error Detection on Computer Performance," 24 pages, October 1982.

abstract -

Conventionally, reliability analyses either assume that a fault/error is detected immediately following its occurrence, or neglect damages caused by latent errors. Though unrealistic, this assumption has been imposed in order to avoid the difficulty of determining the respective probabilities that a fault induces an error and the error is then detected in a random amount of time after its occurrence.

As a remedy for this problem, in this paper a model is proposed to analyze the impact of error detection on computer performance under moderate assumptions. Error latency - the time

interval between occurrence of an error and the moment of error detection - is used to measure the effectiveness of a detection mechanism. We have used this model to (1) predict the probability of producing an unreliable result, and (2) estimate the loss of computation due to fault and/or error.

CRL-TR-3-82.

B.A. Makrucki and T.N. Mudge

"A Stochastic Model of Parallel and Concurrent Program Execution on Multiprocessors," 81 pages, October 1982.

abstract -

This report summarizes a model developed to allow the evaluation of parallel program execution on multiprocessors. The model is intended for MIMD algorithms in which the individual processors are coupled through their programs' interaction with memory. The model is not intended for SIMD algorithms. Specifically, estimates of processor utilization, execution times of programs or subprograms, and memory bandwidth can be obtained from the model. Earlier research has concentrated on the last of these quantities and a body of research, which might be termed "memory interference models", has evolved. The work reported here goes one step further by allowing the programs to be included in the model. Consequently questions about the performance of the processors running the programs can also be answered.

CRL-TR-4-82.

Vaclav Rajlich

"Determinism in Parallel Systems," 10 pages, October 1982.

Withdrawn -

This report deals with determinism in parallel systems, has been published in "Theoretical Computer Science" 26 (1983) 225-231, North-Holland, and is superceeded by that article. Reprints of the article are available from: Professor Vaclav Rajlich, Univ. of Michigan, CCS Dept., 221 Angel Hall, Ann Arbor, MI 48109.

CRL-TR-5-82.

Wan P. Chiang and Toby J. Teorey

"A Method for Database Record Clustering," 25 pages, August 1982.

abstract -

Record clustering is one of the important problems in physical database design. A heuristic method is proposed for identifying the record types in a database that are to be clustered together so that total access time is minimized for the designed database. A paging environment is assumed. The clustering result, while known to be optimal for hierarchical databases, is shown to be at least near optimal for network databases.

CRL-TR-6-82.

Yann-Hang Lee and Kang G. Shin

"Design and Evaluation of a Fault-Tolerant Multiprocessor Using Hardware Recovery Blocks," 50 pages, August 1982.

abstract -

In this paper we consider the design and the evaluation of a fault-tolerant multiprocessor with a rollback recovery mechanism.

The rollback mechanism is based on the hardware recovery block which is a hardware equivalent to the software recovery block. The hardware recovery block is constructed by consecutive state-save operations and several state-save units in every processor and memory module. When a fault is detected, the multiprocessor reconfigures itself to replace the faulty component and then the process originally assigned to the faulty component retreats to one of the previously saved states in order to resume fault-free execution.

Due to random interactions among cooperating processes and also due to asynchrony in the state-savings, the rollback of a process may propagate to others and multiple-step rollbacks may thus become necessary. In the worst case, when all the available saved states are exhausted, the processes have to restart from the beginning as if they were executed in a system without any rollback recovery mechanism. A mathematical model is proposed to calculate both the coverage of multi-step rollback recovery and the risk of restart. The performance evaluation in terms of the mean and variance of execution time of a given task is also presented.

CRL-TR-7-82

Liang Tai Wu

"Models for Evaluating the Performability of Degradable Computing Systems," 193 pages, June 1982.

abstract -

Recent advances in multiprocessor technology have established the need for unified methods to evaluate computing systems performance and reliability. In response to this modeling need, this dissertation considers a general modeling framework that permits the modeling, analysis and evaluation of degradable computing systems. Within this framework, several user-oriented performance variables are identified and shown to be proper generalizations of the traditional notions of system performance and reliability. Furthermore, a time-varying version of the model is

developed to generalize the traditional fault-tree reliability evaluation methods of phased missions.

The modeling and evaluation methods considered in this dissertation provide a relatively straightforward approach to integrate reliability and availability measures with performance measures. The hierarchical decomposition approach permits the modeling and evaluation of a computing system's subsystems (e.g., hardware, software, peripherals, interfaces, user demand systems) as a whole rather than the traditional methods of evaluating these subsystems independently. Accordingly, it becomes possible to evaluate the performance of the system software and the reliability of the system hardware simultaneously in order to measure the effectiveness of the system design. Moreover, since the performance variables considered in this study permit the characterization of system performance according to the application needs of a system, the results obtained represent more accurate assessments of the system's ability to perform than the existing performance of reliability measures.

CRL-TR-8-82

Scott McFarling, Jerry Turney and Trevor Mudge

"VLSI Crossbar Design Version Two," 18 pages, February 1982.

abstract -

This report describes the design of a crossbar switch. A crossbar switch can be used to connect multiple processors to multiple memory modules, or simply as a means of constructing a multiport memory. True concurrent memory accesses are then possible for multiprocessing and DMA. However, memory conflicts can still occur if more than one access is made to the same memory module. In crossbar terms this corresponds to more than one input requesting the same output. In any application in which it is known that input requests can conflict in this way (e.g. if the crossbar is used in an MIMD system) logic circuitry must be provided to resolve possible

conflicts. This logic can become very complex. In the design presented here this problem is simplified by using a technique which we have termed "free for all": the condition where more than one input requests an output is detected and the inputs have to try again at a later date. This requires minimal logic. The retry policy can be determined by whatever is using the inputs. Fabrication of a prototype is presently being undertaken jointly by the ECE department at the University of Michigan and General Motors Technical Center.

CRL-TR-9-82

Andreas Blass and Yuri Gurevich

"Equivalence Relations, Invariants, and Normal Forms," 17 pages, December 1982.

abstract -

For an equivalence relations E on the words in some finite alphabet, we consider the recognition problem (decide whether two words are equivalent), the invariant problem (calculate a function constant on precisely the equivalence classes), the normal form problem (calculate a particular member of an equivalence class, given an arbitrary member), and the first member problem (calculate the first member of an equivalence class, given an arbitrary member). A solution for any of these problems yields solutions for all earlier ones in the list. We show that, for polynomial time recognizable E , the first member problem is always in the class Δ_2^P (solvable in polynomial time with an oracle for an NP set) and can be complete for this class even when the normal form problem is solvable in polynomial time. To distinguish between the other problems in the list, we construct an E whose invariant problem is not solvable in polynomial time with an oracle for E (although the first member problem is in $NP^E \cap co-NP^E$) and we construct an E whose normal form problem is not solvable in polynomial time with an oracle for a certain solution of its invariant problem.

CRL-TR-1-83

Andreas Blass and Yuri Gurevich

"On the Unique Satisfiability Problem," 13 pages, January 1983.

abstract -

UNIQUE SAT is the problem of deciding whether a given Boolean formula has exactly one satisfying truth assignment. We address the question whether UNIQUE SAT is complete for the class $DIF^P = \{L_1 - L_2 : L_1, L_2 \in NP\}$. We construct an oracle relative to which UNIQUE SAT is not complete for DIF^P , and another oracle relative to which UNIQUE SAT is complete for DIF^P whereas $NP \neq co-NP$. (We care about $NP \neq co-NP$ because the equality $NP = co-NP$ easily implies that UNIQUE SAT is complete for DIF^P .)

CRL-TR-2-83

Wen-Tai Liu

"Techniques for Estimation of the Area of Integrated Digital Circuits," 152 pages, January 1983.

abstract -

This dissertation describes a sub-system of an Arithmetic Design System (ADS) which is intended to estimate figures of merit for a particular arithmetic design at the VLSI technology realization level. We emphasize the gate array, the programmable logic array (PLA) and the general cell approach. The fact that the design processes are so time consuming motivates us to develop estimation methods for figures of merit (*i.e.*, area) resulting from design characteristics. These estimates offer advice which avoids wasting design time to achieve only a small amount of improvement. We develop the estimation methods for the gate array and the PLA approaches. In both cases, Rent's relationship represents the design characteristics.

In the gate array approach, we propose a point model for the layout study and show both the final area and the average interconnection length increase as Rent's exponent r increases. In the PLA approach, based on a partition model for the folding process, the saved area ratio δ is shown to be bounded by $\frac{1}{2^r} \leq \delta \leq \frac{1}{2}$.

We propose a constructive layout approach for the general cell model, which is needed to map the hierarchical design description into physical structures. The routability problem and channel router are particularly discussed in detail. The minimum condition set for the routability test and channel routing order generation algorithm are derived by using a graph model to represent all the legal routing orders. The routing process can be entirely completed by sequentially applying a channel router to each channel according to the routing order. The new channel router consists of two graphs, namely the interval graph G_I and the constraint graph G_C , which represent respectively the overlap and constraint relations among the nets. The reduced graph G_r is formed by removing the constraint edges from the complement of G_I . The router chooses the best candidate from the dynamically updated G_r .

CRL-TR-3-83

W.G. Golson and W.C. Rounds

"Connections Between Two Theories of Concurrency: Metric Spaces and Synchronization Trees,"

27 pages, January 1983.

abstract -

We establish a connection between the semantic theories of concurrency and communication in the works of de Bakker and Zucker, who develop a denotational semantics of concurrency using metric spaces instead of complete partial orders, and Milner, who develops an algebraic semantics of communication based upon observational equivalence between processes. We endow his *rigid*

synchronization trees (RSTs) with a simple pseudometric distance induced by Milner's *weak equivalence* relation and show the quotient space to be complete. We establish an isometry between our space and the solution to a domain equation of de Bakker and Zucker, presenting the solution in a conceptually simpler framework. Under an additional assumption, we establish the equivalence of the weak equivalence relation over RSTs and the elementary equivalence relation induced by the sentences of a modal logic due to Hennessy and Milner.

CRL-TR-4-83

C-W. Chung

"A Query Optimization in Distributed Database Systems," 177 pages, February 1983.

abstract -

This research is concerned with a model and a method of minimizing the inter-site data traffic incurred by a query in distributed relational database systems. In order to process a query which references data from multiple sites in a computer network, portions of the database at other sites have to be transferred to the user's site. The usual methodology for distributed query processing consists of reducing the referenced relations using a sequence of semijoin operations after initial local processing.

The mathematical model has been developed to determine an optimal sequence of semijoins which minimizes the total inter-site data flow in processing a distributed query. The core of this model is a method which efficiently and accurately estimates the size of an intermediate result of a query. In particular, the assumption that joining attributes are independent during the processing of a query by a sequence of semijoins has been relaxed.

Since the distributed query optimization problem is known to be NP-hard, a heuristic algorithm has been developed to determine a low-cost sequence of semijoins. The efficiency of the algorithm is increased by partitioning the set of joining attributes into blocks and sequencing

these blocks, as well as by a straightforward, yet effective sequencing among the semijoins between the joining attributes inside a block. The algorithm decreases the cost of a query by selecting the low-cost, highly reductive semijoins first. Cost comparisons with the existing algorithms have been provided. The time complexity of the main features of the algorithm has been analytically derived.

The algorithm has been implemented in PASCAL. The tests show that the scheduling time for a sequence of semijoins for a reasonable size query is less than 0.05 seconds when the program is executed by a main-frame computer.

CRL-TR-5-83

Withdrawn

CRL-TR-6-83

Withdrawn

CRL-TR-7-83

M. Maletz

"On Enhancing the Idealized CPU-I/O and I/O-I/O Overlap Models Through the Use of Markov Processes," 14 pages, April 1983.

abstract -

A Modeling technique is developed for studying processor overlap in time-shared and multiprogrammed computing systems, based on the Markov process methodology. Measures of resource utilization, response time, and system throughput are derived from the statistics of the

Markov process equilibrium. Techniques for identifying potential system bottlenecks and for determining summary statistics about user transition behavior are also discussed. The result is an enhancement of traditional overlap models with greater generality and improved system characterization capabilities.

CRL-TR-8-83

C.M. Krishna and K.G. Shin

"Queueing Analysis of a Canonical Model of Real-time Multiprocessors," 26 pages, February 1983.

abstract -

Multiprocessors are beginning to be regarded increasingly favorably as candidates for controllers in critical real-time control applications such as aircraft. Their considerable tolerance of component failures together with their great potential for high throughput are contributory factors.

In this report, we present first a logical classification of multiprocessor structures from the point of view of control applications. We point out that one important subclass has hitherto been neglected by the analysts. This is the class of systems with a common memory, minimal interprocessor communication and perfect processor symmetry.

The performance characteristic of the greatest importance in real-time applications is the response time distribution. Indeed, we have shown in a separate report [2] how it is possible to characterize rigorously and objectively the performance of a real-time multiprocessor given the application and the multiprocessor response time distribution and component failure characteristics. We therefore present here a computation of the response time distribution for a canonical model of real-time multiprocessor.

To do so, we approximate the multiprocessor by a blocking model and present a means for efficient analysis. Two separate models are derived: one created from the system's point of view, and the other from the point of view of an incoming task. The former model is analyzed along largely conventional lines. For the latter model, an *artificial server* is used and the system transformed into a queueing network. Validations show that the approximation is good.

CRL-TR-9-83

Y-H. Lee and K.G. Shin

"Analysis of Backward Error Recovery for Concurrent Processes with Recovery Blocks," 28 pages, February 1983.

abstract -

Although backward error recovery with recovery blocks(RB's) has received considerable attention from many researchers, no attempt has been made to structure its implementation alternatives and then to evaluate/analyze their effectiveness. In this report we consider three different methods of implementing RB's. These are the asynchronous, synchronous, and the pseudo recovery point implementations.

Asynchronous RB's are based on the concept of maximum autonomy in each of concurrent processes. Consequently, establishment of RB's in a process is made independently of others and unbounded rollback becomes a serious problem.

In order to completely avoid unbounded rollback, it is necessary to synchronize the establishment of recovery blocks in all cooperating processes. Process autonomy is sacrificed and processes are forced to wait for the commitment to establishing a recovery line, leading to inefficiency in time utilization.

As a compromise between asynchronous and synchronous RB's, we propose to insert pseudo recovery points so that unbounded rollback may be avoided while maintaining process autonomy.

We have developed probabilistic models for analyzing these three methods under standard assumptions in computer performance analysis, i.e. exponential distributions for related random variables. With these models we have estimated (i) the interval between two successive recovery lines for asynchronous RB's, (ii) mean loss in computation power for the synchronized method, and (iii) additional overhead and rollback distance in case PRP's are used.

CRL-TR-10-83

R.A. Rutenbar, T.N. Mudge and D.E. Atkins

"A Class of Cellular Architectures to Support Physical Design Automation," 55 pages, February 1983.

abstract -

Special-purpose hardware has been proposed as a solution to several increasingly complex problems in design automation. This paper examines a class of cellular architectures--called *raster pipeline subarrays*-- applicable to physical design automation problems represented on a cellular grid. Machines with this architecture were first employed for cellular image processing, and many similarities exist between problems in grid-based DA and problems in cellular image processing and pattern recognition. A review of machines designed for cellular image processing shows how DA machines proposed/constructed for grid-based problems fit naturally into a taxonomy of image processors; a review of some of the mathematical tools developed to formalize pattern recognition problems shows how they can be usefully applied to DA problems. Implementations of design rule checking and routing algorithms are described in detail for an existing raster pipeline subarray machine called a *cytocomputer*. Experimental results using this hardware are encouraging, and extensions to large, practical problems are studied. Based upon these studies we define the architecture and necessary performance characteristics of a raster pipeline subarray machine optimized specifically for grid-based DA applications. The merits of such an architecture

are evaluated in the context of practical special-purpose hardware.

CRL-TR-11-83

T.N. Mudge and T.A. Rahman

"Efficiency of Feature Dependent Algorithms for the Parallel Processing of Images," 13 pages, March 1983.

abstract -

In this paper the concept of feature (in)dependent image processing algorithms is defined. A large class of image processing computers characterized by multiple processor-memory subsystems is efficient when dealing with feature independent algorithms but less efficient when dealing with feature dependent algorithms. Typically such machines are required to perform both types of algorithms. This paper is a preliminary attempt to provide a framework within which to model feature dependent algorithms, and to, for example, quantify the inefficiency that can occur when they are executed on the above type of parallel image processors.

CRL-TR-12-83

M. Maletz

"Computing System Simulation Using Markov Processes," 13 pages, April 1983.

abstract -

Models of computing systems involving fixed system resources and probabilistically specified user state transition behavior must account for interactions between the system users and the system states that result from resource limitations. This is normally accomplished by the introduction of queues for the system states. Because of the probabilistic definition of user behavior,

queue transition probabilities must usually be determined by a simulation of the computing system. This paper develops a new technique for simulating computing systems where the hardware configuration is specified by a processor/state service matrix and is fixed throughout the simulation, and user state transition behavior is specified using a basic state transition probabilities matrix satisfying the requirements for a Markov process.

CRL-TR-13-83

V. Rajlich

"Stepwise Refinement Revisited," 23 pages, March 1983.

abstract -

In this paper, rigorous application of stepwise refinement is explored. The steps of definition, decomposition, and completion are described, where steps of completion is a newly introduced step. This combination of steps extends the use of stepwise refinement to medium-to-large systems. Notions of range, active objects, and composite interface are introduced. Verification of incomplete programs via interactive testing is described. The paradigm is demonstrated in an example. The relationship between the paradigm and the current programming languages is considered. It is argued that WHILE-DO loop is a harmful construct from this point of view.

CRL-TR-14-83

E. Delp

"The CRC Plotting Package," 44 pages, March 1983.

abstract -

This report describes device independent plotting software that runs on the Computer and Image Processing Research Network at the University of Michigan. This software supports various graphics devices and is relatively portable from one UNIX installation to another.

CRL-TR-15-83

E. Delp

"PDS Users' Manual: Introduction," 99 pages, March 1983.

abstract -

This report describes image processing software running on the Computer and Image Processing Research Network at the University of Michigan (CIPRNET). The software runs under UNIX and is portable.

The package includes various image processing tools and runtime library support. PDS also supports the RAMTEK and DEANZA systems of CIPRNET.

CRL-TR-16-83

K. Hadavi

"A Definitional Technique for Specification and Implementation of Data Types," 179 pages, April 1983.

abstract -

A model is proposed for specification and implementation of data types. This model is based on a novel multi-level graph structure, viz. the Σ -structures. A universal set of data structure operators are defined to characterize the data types. Constructive definitions of these operators are presented to the extent that the task of defining them is reduced to specification of only

three first order predicate expressions. Using these predicates, the important issue of error is easily and automatically taken care of. This is achieved through the use of functions whose domains are defined by the above predicates in such a way that every constructor operation results in a "non-error" configuration.

The type manipulation operations (TMO) are introduced in order to define data types of a different behaviour to those of the existing ones. An example of a TMO is the "embed" operation. It enables one to combine two data types so that the resulting data type exhibits a behaviour which can be automatically derived from the operand data types. This facility may also be used to parameterize data types. Another TMO is the enrichment operation. Two types of enrichment are introduced. These are enrichment for "convenience", and enrichment for "change of behaviour." Sufficient conditions are developed in order to distinguish one from the other.

It is demonstrated that the proposed model is highly "extensible." For example, with a very minor alteration, a stack specification may be changed to that of a queue and vice versa. Also demonstrated is the ability of the specification to define parallelism of the operations. To this end, without any extra burden on the user, highly parallel operations may be defined for updating and/or accessing data.

It is shown that our specification technique offers an invaluable tool to ensure the "security" of a data base, or an operating system. It is also demonstrated that different views of the same set of data may be held by different users concurrently by assigning different predicates to each user.

CRL-TR-17-83

G. Qadah

"A Relational Database Machine: Analysis and Design," 247 pages, April 1983.

abstract -

The collection of data in the form of an integrated database is a sound approach to data management. The conventional implementation of the database system, that is, augmenting a large general-purpose von Neumann computer with a large complex software system, the database management system (DBMS), suffers from several disadvantages. Among these are low reliability and poor performance in support of a large class of database operations. The importance of database systems, the disadvantages of its conventional implementation, the advancement of processor-memory technology, and the continuous drop in its fabrication cost inspired a new approach to database system implementation. This approach replaces the general-purpose von Neumann computer with a dedicated machine, the database machine (DBM), tailored for the data processing environment and, in most cases, utilizing parallel processing to support some or all the functions of the DBMS. The new approach claims to improve database system reliability and performance.

The general framework of this thesis is the design of a DBM suitable for supporting concurrent, on-line, very large relational database systems. In designing this machine, a structured approach is followed. First, the relational data model, together with its most important operations and the previously proposed DBMs, is reviewed. This review, coupled with the requirements of the very large database systems and the restrictions imposed by the current and the anticipated state of technology, is used to formulate a set of design guidelines. Consequently, an architecture for a cost-effective DBM that meets this set of guidelines is synthesized. A review of the previously proposed DBMs is carried out using a novel classification scheme. This scheme not only aids one to understand the various organizations of the previously proposed DBMs as well as their design trade-offs and limitations, but also provides a tool to qualitatively analyze and compare DBM effectiveness in supporting the requirements of the very large database systems.

Within the context of the proposed machine, the implementation of a very important relational algebra operation, the equi-join operation, is extensively studied. A large set of algorithms

is suggested to implement the equi-join operation on the DBM. An average-value modeling technique is proposed and used to evaluate the equi-join implementations and determine the best performing ones.

Finally, the implementation of the other relational algebra operations as well as other primitives essential to the new DBM are developed.

CRL-TR-18-83

T.N. Mudge, G.D. Buzzard, D.J. Verhaeghe, J. Hill, and D.C. Winsor

"Object-based Computer Architectures," 24 pages, April 1983.

abstract -

There has been a sizable increase of interest in object-based computer systems recently. Much of this increase is attributable to DoD's massive commitment to the Ada language project. Though Ada (DoD's proposed standard language) may not fit everybody's definition of an object-based language, it does incorporate key object-based concepts. In this paper we attempt to characterize these systems and the underlying concepts of the object-based design methodology. We also present case studies of two commercially available object-oriented computers and point to issues which require further study.

CRL-TR-19-83

Siamak Arya

"Optimal Instruction Scheduling for a Class of Vector Processors: An Integer Programming Approach," 133 pages, April 1983.

abstract -

An integer programming model that portrays the architectural features of a class of vector and array processors has been developed. This model is used to produce optimal schedules for low-level-instruction codes of such processors. Optimal schedules are produced for both straight codes and instruction loops. The model is extended to optimally reassign registers to instructions in addition to instruction sequencing. The model is further extended to consider processors with multiple identical functional units. A study of the complexity of the model shows that the scheduling time increases exponentially with the number of instructions. Using the model, a number of experiments have been conducted in optimal scheduling of Cray assembly codes.

CRL-TR-20-83

Janice M. Jenkins

"Symposium on Computer Applications to Cardiology Introduction and Automated Electrocardiography and Arrhythmia Monitoring," 47 pages, April 1983.

abstract -

This paper informs us that the electrocardiogram was the first physiologic signal to be processed and analyzed by digital computer, a technique which dates back to 1957. Early attempts at computer-derived measurement of the waveforms progressed to pattern recognition methods for discrimination and classification, and eventually to diagnostic interpretation of the electrocardiogram (ECG). A historical review is presented of the evolutionary stages of computer interpretation of the diagnostic ECG, and the reader is acquainted with various methods employed for signal processing, techniques for waveform and contour analysis, and algorithms for diagnostic interpretation. Commercial systems which evolved from those developed in research settings are described and distinctive features are pointed out.

Computer-based arrhythmia monitors for coronary care were a natural outgrowth of the interpretive system but ineffective P-wave recognition continues to be the major limiting factor in

rhythm analysis. The development of a miniaturized esophageal electrode for computer-detection of P-waves provided a dramatically improved registration of atrial activity and enabled for the first time automated analysis of complex arrhythmias. The two-lead system which recognizes P-waves on the esophageal signal, QRS complexes on the surface lead, and computes PP, RR, and PR intervals is described in Section V of the paper.

Initial attempts at evaluation of the diagnostic ECG have been reported but no consistent scheme has emerged which can be applied in general for comparative purposes. Proposals for the development of a test library with clinical documentation from non-ECG sources have been advanced, but the complexity of the task has served to inhibit its implementation. A promising move in this direction is the creation of an annotated data base for arrhythmia monitors developed under the auspices of the American Heart Association for purposes of testing automated detection systems.

The introduction of computers into clinical electrocardiography has not resulted in any widespread improvement of diagnostic accuracy or dramatically altered the delivery of medical care, but the technique of computer analysis and storage of large numbers of ECGs within a single system provides a powerful tool for epidemiologic studies and possible early detection of cardiovascular disease. Technologic advances include remote acquisition and telephonic transmission of ECGs to a central computer facility and advanced electrocardiographs with internally contained microprocessors which instantly deliver a diagnostic report at the patient site.

CRL-TR-21-83

Yuri Gurevich

"Algebras of Feasible Functions," 10 pages, May 1983.

abstract -

For different complexity levels (PTIME, LOGSPACE, etc.) and an arbitrary set σ of functions we give an inductive definition of the class of functions computable from σ within the complexity level. The idea is to view computable functions as database queries rather than pure arithmetical objects. For example, the function

$$f^G(x) : \text{the diameter of the connected component of vertex } x \text{ of graph } G$$

can be coded into a pure arithmetical function. We prefer, however, to view it as a kind of global function (or a function schema) that becomes an ordinary function in each graph. Global functions are defined precisely in §1. The usual definition of primitive recursive functions, adapted to global functions, surprisingly gives exactly LOGSPACE computable global functions, see §2. Recursive global functions appear to be exactly PTIME computable global functions, see §3. To show that our approach can handle some other complexity classes, we mention some more results in §4.

CRL-TR-22-83

C.S.G. Lee and B.H. Lee

"Resolved Motion Adaptive Control for Mechanical Manipulators," 23 pages, June 1983.

abstract -

This paper presents the development of a resolved motion adaptive control which adopts the ideas of "Resolved Motion Rate Control" [9] and "Resolved Motion Acceleration Control" [10] to control a manipulator in Cartesian coordinates for various loading conditions. The proposed adaptive control is performed at the hand level and is based on the linearized perturbation system along a desired hand trajectory. A recursive least square identification scheme is used to perform on-line parameter identification of the linearized perturbation system. The controlled system is characterized by feedforward and feedback components which can be computed separately and

simultaneously. The feedforward component resolves the specified positions, velocities, and accelerations of the hand into a set of values of joint positions, velocities, and accelerations from which the nominal joint torques are computed using the Newton-Euler equations of motion to compensate all the interaction forces among the various joints. The feedback component computes the variational joint torques which reduce the manipulator hand position and velocity errors along the nominal hand trajectory. This adaptive control strategy reduces the manipulator control problem from a nonlinear control to controlling a linear control system about a desired hand trajectory. The feasibility of implementing the proposed adaptive control using present day low-cost microprocessors is discussed.

CRL-TR-23-83

K.G. Shin, C.M. Krishna and Y.H. Lee

"A Unified Method for Evaluating Real-time Computer Controllers: A Case Study," 45 pages, June 1983.

abstract -

A real-time control system consists of a synergistic pair, that is, a *controlled process* and a *controller computer*. We have defined new performance measures for real-time controller computers on the basis of the nature of this synergistic pair.

In this report we present a case study of a typical critical controlled process in the context of new performance measures that express the performance of both controlled processes and real-time controllers (taken as a unit) on the basis of a single variable: *controller response time*. Controller response time is a function of current system state, system failure rate, electrical and/or magnetic interference, etc., and is therefore a random variable. Control overhead is expressed as a monotonically non-decreasing function of the response time and the system suffers catastrophic failure, or *dynamic failure*, if the response time for a control task exceeds the corresponding

system *hard deadline*, if any. A rigorous probabilistic approach is used to estimate the performance measures.

The controlled process chosen for study is an aircraft in the final stages of descent, just prior to landing. Control constraints are particularly severe during this period, and great care must be taken in the design of controllers that handle this process. First, the performance measures for the controller are presented. Secondly, control algorithms for solving the landing problem are discussed and finally the impact of our performance measures on the problem is analyzed, showing that the performance measures and the associated estimation method have great potential use for designing and/or evaluating real-time controllers and controlled processes. Also, one application for the design of controller computers, presented in detail, is checkpointing for enhanced reliability.

CRL-TR-24-83

Michael R. Betker, E.J. Delp and T.N. Mudge

"An Investigation of the Discrete Karhunen-Loeve Transform: Methods and Computational Aspects," 17 pages, June 1983.

abstract -

In this report we discuss some of the computational aspects of the discrete Karhunen-Loeve transform. We examine the implementation of the algorithm in image data compression.

CRL-TR-25-83

William C. Rounds

"On the relationships between Scott domains, synchronization trees, and metric spaces," 2 pages, July 1983.

abstract -

Scott's theory of Information Systems is intended to provide an easy way to define partial order structures (domains) for denotational semantics. This paper illustrates the new method by considering a simple modal logic, due to Hennessy and Milner, as an example of an information system. The models of formulas in this logic are the rigid synchronization trees of Milner. We characterize the domain defined by the Hennessy-Milner information system as the complete partial order of *synchronization forests*: nonempty closed sets of synchronization trees. "Closed" means closed with respect to a natural metric distance on synchronization trees, first defined by de Bakker and Zucker and characterized by Golson and Rounds. The metric space methods are used to extend natural tree operations to forests. These operations become sup-continuous when extended, and therefore can be used to provide a denotational semantics for concurrency which allows the full power of least fixpoint methods for recursion.

CRL-TR-26-83

B.A. Makrucki and T.N. Mudge

"A Queueing Model of Delta Networks," 25 pages, August 1983.

abstract -

This report describes an analysis of multistage interconnection networks where queues are placed in the $b \times b$ crossbar switches on which the networks are based. A queueing analysis of the network is presented, and results are obtained using approximations that are appropriate for network operation parameters of primary interest. From the analysis communication delay time and network throughput are derived. Using the results obtained, queue lengths may be chosen so that the network satisfies certain performance requirements.

CRL-TR-27-83

D.G. Furchtgott and J.F. Meyer

"Performability Evaluation of Computing Systems Using Reward Models," 23 pages, August 1983.

abstract -

An evaluation algorithm is developed for a broad class of performability models wherein system performance is identified with "reward." More precisely, for a system S and a utilization period T , the performance variable of the model is the reward derived from using S during T . The state behavior of S is represented by a finite-state stochastic process (the base model); reward is determined by reward rates associated with the states of the base model. It is assumed that the corresponding reward model is a nonrecoverable process in the sense that a future state (reward rate) of the model cannot be greater than the present state. For this model class, we obtain a general method for determining the probability distribution function of the performance (reward) variable and, hence, the performability of the corresponding system. Moreover, this is done for bounded utilization periods, and assumption which demands a relatively complex solution. The result is an integral expression which can be solved either analytically or numerically. A program written for numerical solutions is discussed and an example illustrating the method is presented.

CRL-TR-28-83

J.F. Meyer

"Performability Modeling of Distributed Real-Time Systems," 18 pages, August 1983.

abstract -

This survey/position paper concerns modeling concepts and techniques for the performability (performance-reliability) evaluation of distributed real-time systems. Due to the nature of application requirements, such systems typically exhibit properties of concurrency, timeliness, fault tolerance, and degradable performance. Relevant prior research is surveyed and classified according to these properties and, based on the position that performability models should accommodate all four properties, directions for future research are suggested.

CRL-TR-29-83

G.D. Buzzard and T.N. Mudge

"Object-Based Computer Systems and the Ada Programming Language," 34 pages, August 1983.

abstract -

This paper gives an introduction to object-based computer systems. In particular, it is shown how they can support the Ada programming language. A case study using the Intel iAPX432 is given.

CRL-TR-30-83

M.D. Diamond

"The Graph Labeling Model and Its Application To The Problem of Edge Linking," 148 pages, August 1983.

abstract -

This report covers various aspects of the discrete and continuous graph labeling problem and its application to the problem of edge linking in computer vision. Included are results on the convergence time of the discrete relaxation labeling processes, and results pertaining to the 0-1

integer programming formulation of the continuous graph labeling problem. Several heuristics iterative parallel processes for solving the graph labeling problem are processes. One approach is based on a Lagrange Dual formulation, another is based on dynamic programming. The graph labeling model is applied to the problem of linking edges to extract outlines and several examples of this application are given.

CRL-TR-31-83

V. Rajlich

"A Paradigm for Top-Down Design with Packages" 48 pages, November 1983.

abstract -

The paper presents a paradigm for the design with packages. The paradigm combines the methodology of stepwise refinement with information hiding principle. It is based on steps of definition, decomposition, completion, and abstraction, which produce clean and well-defined packages. A part of the paradigm are the rules by which the procedures acquire the parameters (called first and second rules for parameters). The paradigm is illustrated by an example.

It is argued that "with" clauses of ADA (together with ADA order of compilation) encourage bottom-up programming while discouraging the top-down programming, and hence should be considered harmful.

CRL-TR-32-83

C.M. Krishna, K.G. Shin and R.W. Butler

"Synchronization and Fault-Masking in Redundant Real-Time Systems" 48 pages, November 1983.

abstract -

A real-time computer may fail because of (i) massive component failures or (ii) not responding quickly enough to satisfy real-time requirements. An increase in redundancy -- a conventional means of improving reliability -- can improve the former but can -- in some cases -- degrade the latter considerably due to the overhead associated with redundancy management, namely the time delay resulting from synchronization and voting/interactive consistency techniques.

In this report, we consider the implications of synchronization and voting/interactive consistency algorithms in N-modular clusters on reliability. All these studies have been carried out in the context of real-time applications. As a demonstrative example, we have analyzed results from experiments conducted at the NASA Airlab on the Software Implemented Fault-Tolerance (SIFT) computer. This analysis has indeed indicated that in most real-time applications, it is better to employ hardware synchronization instead of software synchronization, and not allow reconfiguration.

CRL-TR-33-83

K.G. Shin and C.M. Krishna

"New Performance Measures for Design and Evaluation of Real-Time Multiprocessors" 48 pages, November 1983.

abstract -

Conventional performance measures such as throughput, reliability, and availability are not by themselves alone suitable for analyzing and designing real-time multiprocessors. (The real-time system we focus on here consists of a *computer controller and controlled processes*). Most of these measures treat the computer as a monolith in terms of the services it delivers, thus failing to accurately describe the nature of diversified real-time application tasks.

As a remedy for this problem, this report presents some new performance measures based on computer(multiprocessor) controller response time that are (i) *congruent* to the real-time applications, (ii) able to offer an *objective comparison* of rival computer systems, and (iii) experimentally *measurable/determinable*. These measures, unlike others, provide the real-time multiprocessor controller with a *natural link* to controlled processes.

In order to demonstrate their utility and power, these measures are first determined for an example controlled process on the basis of two control performance functionals. They are then used for an important real-time multiprocessor design application -- the number-power tradeoff. Also, some potential applications of these measures are identified.

CRL-TR-34-83

M.A. Shah and R.C. Jain "Detecting Time-Varying Corners"

abstract -

The algorithms for *structure from motion* require solution of the correspondence problem. By detecting only time-varying tokens, the problem may be significantly simplified. In this paper, a time varying corner detector is described which is based on the **and** operation between the **cornerness** and the temporal derivative. It is shown that the corner detectors by Zuniga and Haralick, Kitchen and Rosenfeld, and Dreschler and Nagel are equivalent. In our time-varying corner detector, we use Zuniga and Haralick corner detector for finding the **cornerness** at a point and use absolute value of difference in intensity at a point to approximate the temporal derivative. The results of the time varying corner detector for the real scenes and the synthetic images with random background and random object are shown.

CRL-TR-35-83

Michael R. Betker, E.J. Delp and T.N. Mudge "An Investigation of the Discrete Karhunen-Loeve Transform: Methods and Computational Aspects"

abstract -

In this report we discuss some of the computational aspects of the discrete Karhunen-Loeve transform. We examine the implementation of the algorithm in image data compression scheme. We further discuss extension of the algorithm using a high speed floating point computer system.

UNIVERSITY OF MICHIGAN



3 9015 02653 6055