Additional Papers

In the December 1991 issue of *The Computer Journal* a new concept of *Additional Papers* was introduced. Such papers will have been typeset and given numbered pages in the same way as other papers, but only the title, authorship and abstract appear in the *Journal*. To obtain copies of the full paper, readers apply to the editor who will mail copies directly to the reader at cost. Libraries will be able to obtain bound volumes at the year end. Reprints will be supplied to authors. For further information, including price lists and order forms, readers and libraries should apply directly to the Editor-in-Chief at the address given on the inside front cover of the *Journal*.

A further list of additional papers is given below.

Dynamic File Organizations For Partial Match Retrieval Based On Linear Hashing

T. S. YUEN AND H. C. DU, *Computer Science Department, University of Minnesota, 4-192 EE/CSci Building, 200 Union Street, S.E., Minneapolis, MN55455-0159, USA*

Two new file organizations based on linear hashing are proposed for partial match retrieval. The first organization introduces a load-balancing scheme whereby overflow records are stored temporarily in other primary buckets so that the allocation of overflow buckets are deferred. The second organization defers the physical splitting of underflow buckets, so that the records belonging to underflow buckets can be retrieved together. These two techniques are then combined together to form a new variation of linear hashing. Compared with the original scheme, the performance of these organizations for partial match retrieval are improved, both in terms of storage utilization and retrieval time.

(pp. A467–A474)

Structuring Spatially Related Meta-Information for Effective Entry and Retrieval

C. L. N. RUGGLES[1], I. A. NEWMAN[2], D. MEDYCKYJ-SCOTT[2] AND D. R. F. WALKER[3], [1]*Midlands Regional Research Laboratory and Department of Computing Studies, University of Leicester,* [2]*Midlands Regional Research Laboratory and Department of Computer Studies, Loughborough University of Technology, Loughborough, Leicestershire LE11 3TU,* [3]*Midlands Regional Research Laboratory and Department of Geography, Loughborough University of Technology*

This paper discusses some aspects of a general information structuring problem. The particular context which is considered relates to the entry and retrieval of meta-information (i.e. information about information) by users of the services provided by the Midlands Regional Research Laboratory (MRRL). The concept is one of providing access to a large number of datasets relevant to a region, in a variety of forms and formats, including a central database of its own. 'Providing access to' such datasets involves holding knowledge of their structure in order to allow a user to formulate queries spanning one or more of them. An attempt is made in this paper to develop an abstract model for structuring unconstrained, spatially related information in order to provide effective entry and retrieval, with the particular requirements of the MRRL metadatabase providing the primary motivation. The development is formal, and is expressed using VDM specification language.

(pp. A475–A486)

A Controlled Environment for Semantic Information Processing within Inter-process Communication

E. J. YANNAKOUDAKIS[1] AND A. A. SHAHID[2], [1]*Department of Computer Science, University of Bradford, Bradford BD7 1DP, W. Yorkshire,* [2]*National Institute of Electronics, P.O. Box 1406, H-9, Islamabad, Pakistan*

This paper deals with dynamic information concerning the activity of a computer system. This information is vital to enable the user to monitor the system's behaviour and other processes to synchronise and cooperate with each other. The paper gives a new approach to semantic information extraction by concurrent execution monitoring and its presentation, in a form interpretable by another process and understandable by the user. An automated mechanism for inter-process communication and system monitoring is also presented.

(pp. A487–A497)

Hash Functions for Hash-based Join Methods

J. K. MULLIN, *Computer Science Department, University of Western Ontario, London, Ontario, Canada N6A 5B7*

This paper discusses two issues in hash-based joins in database systems:

The optimal number of split files [a split is a file used to partition a larger file so that the smaller sections can be conveniently joined] for the join: This choice is roughly analogous to choice of the order of merge in sorting. It determines the range of the hash functions.

Appropriate hash functions to achieve this split: An appropriate choice of split has an important effect on join algorithms – these are known to be a major item affecting cost in relational database systems. It will be shown that some conventional methods such as division remainder hashing and the Universal hash functions of Carter and Wegman are inappropriate in this application. Another method, which is a slight modification of Universal hash functions, will be shown to work well with alphabetic data.

(pp. A499–A503)

A Generic Adaptive Sorting Algorithm

V. ESTIVILL-CASTRO AND D. WOOD, *Faculty of Mathematics, Department of Computer Science, University of Waterloo, Waterloo, Ontario N2L 3G1, Canada*

We present a generic adaptive sorting algorithm that enables us to focus attention on the combinatorial properties of measures of presortedness rather than the combinatorial properties of the algorithm. Using it we obtain a practical adaptive sorting algorithm, optimal with respect to three important measures of presortedness and smoothly adaptive for other common measures. The new sorting algorithm can be easily modified to achieve optimality for two more important measures.

Moreover, we extend the proof techniques to analyse an adaptive variant of* Quicksort*; previous claims were based only on simulation results.

(pp. A505–A512)

Improving Firing Control via Explicit Enabling Rules in Petri Nets

R. H. GALLARD, *Facultad de Ciencias Fisico, Matematicas y Naturales, Universidad Nacional de San Luis, Loçal 106, Ejrecito de los Andes 950 – lo. P. – 5700 San Luis, Republica Argentina*

When modelling partial aspects of operating systems, basic Petri nets have proved useful in studying concurrency. However, they are insufficient to determine some performance measures of the system under design. This problem arises, essentially, due to the following limitations; tokens are not distinguishable from one another and transition firings are ungovernable. This paper introduces controlled Petri nets as an extension to overcome those limitations. The extension allows the insertion of entity attributes and provides different, explicit, enabling rules for transition. While the natural expressiveness of Petri nets is preserved, they become suitable to interact with external environments. An example is given to show the flexibility and usefulness of controlled Petri nets to model the behaviour of a simple process manager.

(pp. A513–A517)

An Efficient Strategy for the Bottom-up Evaluation of Datalog Queries

N. LEONE AND P. RULLO, *CRAI, 87036 Rende, Italy*

Rewriting methods, such as those based on magic sets, are used to improve the efficiency of the bottom-up evaluation of Datalog queries. The basic idea behind such methods is that of "pushing down" possible constants appearing in query goals, thus restricting the computation of the query answer to only relevant facts. However, it may happen that the rewriting of a stratified Datalog program generates an unstratified program, ie a program where negation occurs "within" recursion. Under such a condition, the convergence of the bottom-up evaluation process is not guaranteed any more.

In this paper we present an algorithm for the efficient bottom-up computation of the stratified semantics of Datalog queries. This algorithm, called the SNQS Algorithm, is based on a combination of the semi-naive method with the query/subquery strategy. Actually it is a simple variant of the "pure" bottom-up approach but it has a larger applicability domain. Indeed, it allows the evaluation of unstratified programs deriving from stratified ones when rewriting techniques are applied, thus supporting a full exploitation of query bindings.

(pp. A519–A527)

Errata

It has been drawn to the attention of the Editors that there were a large number of mistakes in the paper 'On the Meaning of Safety and Security' by A. Burns, J. McDermid and J. Dobson, **35** (1), 3–15.

On page 4, column 2, line 4
'out' *should be* 'our'.
On page 4, column 2, section 2.1, lines 16–19
• 'a *safety critical* system is one whose failure could do us immediate, direct harm;
• a *security critical* system is one whose failure could enable, or increase the ability of, others to harm us.'
should be
• 'a **safe** system is one that cannot do us immediate, direct harm *even if it fails*;
• a **secure** system is one that cannot enable others to harm us *even if it fails*.'
On page 4, column 2, lines 22 and 25
'is' (at the end of the lines) *should be* 'us'.
On page 6, column 2, line 29
'when in' *should be* 'when'.
On page 6, column 2, line 30
'upon' *should be* 'open'.
On page 7, column 1, line 31
'even' *should be* 'event'.
On page 7, column 1, line 51
'the' *should be* 'that'.
On page 7, column 1, line 55
'even' *should be* 'event'.
On page 8, column 1, lines 14–15
'these being either ordinary components in an error state is detected.*'
should be
'these being either ordinary components in an error state or components only invoked when an error state is detected.*'.
On page 8, column 1, line 17
'or or' *should be* 'or'.
On page 8, column 2, line 13
'all' *should be omitted.*
On page 8, column 2, line 50
'object' *should be* 'objects'.
On page 8, column 2, line 54
'the' *should be omitted.*
On page 11, column 2, line 24
'to' *should be* 'is'.