

Lecture Notes in Computer Science

Edited by G. Goos and J. Hartmanis

317

Timo Lepistö Arto Salomaa (Eds.)

Automata, Languages and Programming

15th International Colloquium
Tampere, Finland, July 1988
Proceedings



Springer-Verlag

TP31-53
L939 X A939
1988 1988

8963526

Lecture Notes in Computer Science

Edited by G. Goos and J. Hartmanis

317

Timo Lepistö Arto Salomaa (Eds.)



Automata, Languages and Programming

15th International Colloquium
Tampere, Finland, July 11–15, 1988
Proceedings



Springer-Verlag

Berlin Heidelberg New York London Paris Tokyo



E8963526


Editorial Board

D. Barstow W. Brauer P. Brinch Hansen D. Gries D. Luckham
C. Moler A. Pnueli G. Seegmüller J. Stoer N. Wirth

Editors

Timo Lepistö
Tampere University of Technology
P.O. Box 527, SF-33101 Tampere, Finland

Arto Salomaa
Department of Mathematics, University of Turku
SF-20 500 Turku 50, Finland



CR Subject Classification (1987): C.1, E.1, F, G.2, I.1, I.2.3

ISBN 3-540-19488-6 Springer-Verlag Berlin Heidelberg New York
ISBN 0-387-19488-6 Springer-Verlag New York Berlin Heidelberg

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in other ways, and storage in data banks. Duplication of this publication or parts thereof is only permitted under the provisions of the German Copyright Law of September 9, 1965, in its version of June 24, 1985, and a copyright fee must always be paid. Violations fall under the prosecution act of the German Copyright Law.

© Springer-Verlag Berlin Heidelberg 1988
Printed in Germany

Printing and binding: Druckhaus Beltz, Hemsbach/Bergstr.
2145/3140-543210

PREFACE

ICALP 88 is the 15th International Colloquium on Automata, Languages and Programming in a series of meetings sponsored by the European Association for Theoretical Computer Science (EATCS). It is a broadly based conference covering all aspects of Theoretical Computer Science including topics like Computability, Automata, Formal Languages, Analysis of Algorithms, Computational Complexity, Data Types and Data Structures, Theory of Data Bases and Knowledge Bases, Semantics of Programming Languages, Program Specification, Transformation and Verification, Foundations of Logic Programming, Theory of Logical Design and Layout, Parallel and Distributed Computation, Theory of Concurrency, Symbolic and Algebraic Computation, Term Rewriting Systems, Cryptography, Theory of Robotics.

ICALP 88 was held at the campus of the Tampere University of Technology from July 11 to July 15, 1988. The organizing committee consisted of T. Lepistö, P. Järvinen, R. Kurki-Suonio, K. Ruohonen, K.-J. Räihä, M. Tienari, T. Granroth, P. Kaila and E. Hirvonen.

The program committee selected 46 papers from 195 extended abstracts and draft papers submitted. Each submitted paper was evaluated by at least four members of the program committee. The final selection was made during a two-day selection meeting in which the following program committee members participated: S. Abramsky, G. Ausiello, J. Díaz, J. Gruska, W. Kuich, J. van Leeuwen, M. Nivat, J. Paredaens, A. Paz, A. Salomaa, S. Skyum, E. Ukkonen and D. Wood. It is a pleasure to thank all those who have submitted papers for consideration, the members of the program committee for their help in the evaluation of the papers and the many who assisted in this process.

We also gratefully acknowledge all the institutions and corporations which supported this conference.

Finally, we would like to thank K.-J. Räihä, K. Ruohonen and P. Kaila, who did a beautiful job in all organizational matters related to the conference.

July 1988

Timo Lepistö
Chairman of the
Organizing Committee

Arto Salomaa
Chairman of the
Program Committee

Program Committee

S. Abramsky
A.V. Aho
G. Ausiello
M. Broy
J. Díaz
J. Gruska
G. Huet
W. Kuich
J. van Leeuwen
M. Nivat
Th. Ottman
C.H. Papadimitriou
J. Paredaens
A. Paz
A. Salomaa
S. Skyum
E. Ukkonen
D. Wood

Sponsors of ICALP 88

The main sponsor is IBM, Finland
Other sponsors:
Academy of Finland
Administrative Unit of Häme
City of Pori
City of Tampere
Digital Equipment Corporation
The Ministry of Education,
Finland
Nevanlinna Institute
Nokia Oy Ab
Springer-Verlag

EATCS Institutional Sponsors

Bull, Louveciennes, France
Dansk Datamatik Center/
DDC International, Lyngby
Denmark
Nixdorf Computer AG, Paderborn,
Federal Republic of Germany
Philips Research Lab. Eindhoven,
The Netherlands
Siemens ZTI, Munich, West Germany

List of referees for ICALP 88

Abrial,J.R.	Curien,P.-L.	Heinz,A.
Afonte,M.V.	Damgaard,I.	Hennicker
Astesiano,E.	Ben David,S.	Hesselink,W.H.
Asveld,P.R.J.	Despeccoux,J.	Hofri,M.
Baaz,M.	Duris,P.	Holmström,L.L.
Back,R.-J.	Dybjer,P.	Honkala,J.
de Baer,D.	Ehler	van Horebeek,I.
Baeten,J.C.M.	Eklund,P.	Hromkovic,J.
de Bakker,J.W.	van Emde Boas,P.	Hussmann
Balcázar,J.L.	Engelfriet,J.	Icking,Chr.
Bar Yehuda,R.	Etzion,T.	Indermark,K.
Baron,G.	Farinas del Cerro,L.	Itai,A.
Barringer,H.	Feder,T.	Italiano,G.F.
Bentley,J.	Felgentreu, K.-U.	Janssens,D.
Berry,D.	Floréen,P.	Johnson,D.S.
Berry,G.	Francez,N.	Johnson,J.H.
Berstel,J.	Frandsen,G.	Jouannaud,J.P.
Bertin,P.	Freitag,B.J.	Kahn,G.
Bertolazzi,P.	Frutos Escrig,D.	Kaminski,M.
Bodlaender,H.L.	Gabarro,J.	Kari,J.
Bougé,L.	Gambosi,G.	Karvi,T.
de Bra,P.	Ganzinger,H.	Katz,S.M.
Brell	Gerth,R.	Kelemenová,A.
Bruynooghe,M.	Geser	Klein,R.
Brüggemann-Klein,A.	van Gestel,E.	Klop,J.
Buss,J.F.	Goldreich,O.	Kolaitis,P.G.
Casas,R.	Grahne,G.	Korach,E.
Chan,E.P.F.	Grass,W.	Korec,I.
Cherkasova,L.	Guessarian,I.	Koskimies,K.
Chytil,M.	Guibas,L.J.	Koymans,R.
Codognet,P.	Gyssens,M.	Kranakis,E.
Cohen,E.	Haddad,R.	Kröger,F.
Cohn,E.R.	Hájek,P.	Kuiper,R.
Colbourn,C.J.	Hankin,C.	Küster,G.
Coquand,Th.	Hardin,Th.	Lafont,Y.
Costa,R.	Harju,T.	Larsen,K.G.
Courcelle,B.	Harrison,P.	Laville,A.
Cucker,F.	van Hee,K.M.	Le Chenadec,Ph.

- Lee,D.
 Lenstra,J.K.
 Lescanne,P.
 Levy,J.J.
 Lewi,J.
 Linna,M.
 Lippe,W.M.
 Louchard,S.
 Lubiw,A.
 MacQueen,D.
 Mäkinen,E.
 Mannila,H.
 Marchetti-Speccamela,A.
 Martin,U.
 Mascari,G.
 Mauny,M.
 Mayoh,B.
 Mayr,E.
 Meyer,J.-J.Ch.
 Möncke,U.
 Moran,S.
 Morlock,M.
 Mosses,P.D.
 Nagl,M.
 Nardelli,E.
 Nickl
 de Nicola,R.
 Nielsen,M.
 Niemi,V.
 Nijholt,A.
 Nurmi,O.
 Olderog,E.-R.
 Olivié,H.
 Ong,C.-H.
 Orejas,F.
 Orponen,P.
 Overmars,M.
 Pahl,J.K.
 Padawitz,P.
 Penttonen,M.
 Pepper,P.
 Phillips,I.
 Plambeck,T.
 Poigné,A.
 Privara,I.
 Protasi,M.
 Räihä,K.-J.
 Rathmann,P.K.
 Reisig,W.
 Rem,M.
 Richier,J.-L.
 Ringwood,G.
 Rodriguez Artalejo,M.
 de Roever,W.P.
 Rosenberg,A.
 Ron,E.
 Rozenberg,G.
 Rován,B.
 Roy,S.
 Ruohonen,K.
 Rusinowitch,M.
 Ruzicka,P.
 Ryslinková,J.
 Sadler,M.
 Salomaa,K.
 Saltor,F.
 Sassano,A.
 Schäffer,A.
 Schmerl
 Schmidt,E.M.
 Schmidt,G.
 Schmidt,U.
 Schöning,U.
 Schor,B.
 Schwartzbach,M.
 Shmueli,O.
 Sifakis,J.
 Silberman,G.M.
 de Simone,R.
 Smolka,G.
 Smyth,M.
 van de Snepscheut,J.L.A.
 Soisalon-Soininen,E.
 Soria,M.
 Staunstrup,J.
 Steinby,M.
 Steyaert,J.-M.
 Streicher
 Sturc,J.
 Sturtivant,C.
 Subramanian,A.
 Sykora,O.
 Talamo,M.
 Tamminen,M.
 Taubner,D.
 Thomsen,B.
 Tirri,H.R.
 Tomasta,P.
 Torán,J.
 Torras,J.
 Traktenbrot
 Urbanek,F.J.
 Vaandrager,F.
 Vajtersic,M.
 Vardi,M.Y.
 Vavasis,S.A.
 Veldhorst,M.
 Vickers,S.
 Vitanyi,P.M.B.
 Volger
 Vornberger,O.
 Vyskoc,J.
 van der Weide,Th.P.
 Widmayer,P.
 Wiederman,J.
 Wilhelm,R.
 Wirsing,M.
 Yannakakis,M.
 Yu,S.
 Zaks,S.
 Zucker,J.

TABLE OF CONTENTS

Aggarwal A., Chandra A.K. Communication complexity of PRAMs (Preliminary Version)	1
Albert L., Fages F. Average Case Complexity Analysis of the RETE multi-pattern match algorithm	18
Arnborg S., Lagergren J., Seese D. Problems easy for tree-decomposable graphs (Extended Abstract)	38
Back R.J.R. , Kurki-Suonio R. Serializability in distributed systems with handshaking	52
Bajaj C., Kim M.-S. Algorithms for planar geometric models	67
Benedek G.M., Itai A. Nonuniform learnability	82
Berstel J., Reutenauer C. Zeta functions of recognizable languages	93
Bodlaender H.L. Dynamic programming on graphs with bounded treewidth	105
Chang J.H., Ibarra O.H., Palis M.A. Efficient simulations of simple models of parallel computation by time-bounded ATM's and space-bounded TM's	119
Cole R., Salowe J., Steiger W.L., Szemerédi E. Optimal slope selection	133

Cori R., Métivier Y.	
Approximation of a trace, asynchronous automata and the ordering of events in a distributed system	147
Culik II K.	
New techniques for proving the decidability of equivalence problems (Invited lecture)	162
Diekert V.	
Transitive orientations, Möbius functions, and complete semi-Thue systems for free partially commutative monoids	176
Dietzfelbinger M., Maass W.	
The complexity of matrix transposition on one-tape off-line Turing machines with output tape	188
Edelsbrunner H.	
Geometric structures in computational geometry (Invited lecture)	201
Edelsbrunner H., Guibas L., Pach J., Pollack R., Seidel R., Sharir M.	
Arrangements of curves in the plane - topology, combinatorics, and algorithms	214
Eppstein D.	
Reset sequences for finite automata with application to design of parts orienters	230
Flajolet P., Gardy D., Thimonier L.	
Random allocations and probabilistic languages	239
Gruska J.	
Systolic architectures, systems and computations (Invited lecture)	254
Hartmanis J.	
New developments in structural complexity theory (Invited lecture)	271

Kirchner C., Kirchner H., Meseguer J. Operational semantics of OBJ-3	287
Kirschenhofer P., Prodinger H., Szpankowski W. Do we really need to balance Patricia tries (Extended Abstract)	302
Kok J.N., Rutten J.J.M.M. Contractions in comparing concurrency semantics	317
Kruskal C.P., Rudolph L., Snir M. A complexity theory of efficient parallel algorithms (Extended Abstract)	333
Kucera L., Marchetti-Spaccamela A., Protasi M. On the learnability of DNF formulae	347
Lautemann C. Efficient algorithms on context-free graph languages	362
Lengauer T., Wanke E. Efficient analysis of graph properties on context-free graph languages (Extended Abstract)	379
Lingas A., Syslo M.M. A Polynomial-time algorithm for subgraph isomorphism of two-connected series-parallel graphs	394
Mehlhorn K., Yap C.-K. Constructive Hopf's theorem: or how to untangle closed planar curves	410
Niemi V. Maximal dense intervals of grammar forms	424
Panangaden P., Stark E.W. Computations, residuals, and the power of indeterminacy	439

Rajasekaran S., Reif J.H.	
Nested annealing: A provable improvement to simulated annealing	455
Ramesh R., Ramakrishnan I.V.	
Nonlinear pattern matching in trees	473
Renji Tao	
Invertibility of linear finite automata over a ring	489
Rohnert H.	
Moving discs between polygons	502
Rudich S., Berman L.	
Optimal circuits and transitive automorphism groups	516
Schäfer J.	
A Kleene-Presburgerian approach to linear production systems	525
Simon K.	
On minimum flow and transitive reduction	535
Spehner J.-C.	
La reconnaissance des facteurs d'un langage fini dans un texte en temps linéaire	547
Straubing H., Thérien D., Thomas W.	
Regular languages defined with generalized quantifiers	561
Tamassia R.	
A dynamic data structure for planar graph embedding (Extended Abstract)	576
Tang S., Book R.V.	
Separating polynomial-time Turing and truth-table reductions by tally sets (Preliminary Report)	591
Tel G.	
Assertional verification of a timer-based protocol	600

Thatte S.	
Type inference with partial types	615
Thiagarajan P.S.	
Some behavioural aspects of net theory (Invited lecture)	630
Turakainen P.	
The equivalence of DGSM replications on \mathbb{Q} -rational languages is decidable	654
Vazirani V.V., Yannakakis M.	
Pfaffian orientations, 0/1 permanents, and even cycles in directed graphs	667
Wagner K.W.	
On restricting the access to an NP-oracle	682
Watanabe O.	
On \leq_{1-H}^P Sparseness and nondeterministic complexity classes (Extended Abstract)	697
Weijland W.P.	
Semantics for logic programs without occur check	710
You J.-H.	
Outer narrowing for equational theories based on constructors	727

Communication Complexity of PRAMs

(Preliminary Version)

Alok Aggarwal
Ashok K. Chandra

IBM Research Division
Thomas J. Watson Research Center
P. O. Box 218, Yorktown Heights, New York 10598

Abstract: We propose a model for the concurrent read exclusive write PRAM that captures its communication and computational requirements. For this model, we present several results, including the following:

Two $n \times n$ matrices can be multiplied in $O(n^3/p)$ computation time and $O(n^2/p^{2/3})$ communication delay using p processors (for $p \leq n^3/\log^{3/2}n$). Furthermore, these bounds are optimal for arithmetic on semirings (using $+$, \times only). For sorting and for FFT graphs, it is shown that communication delay of $\Omega(n \log n/(p \log(n/p)))$ is required for $p \leq n/\log n$. This bound is tight for FFT graphs; it is also shown to be tight for sorting provided $p \leq n^{1-\epsilon}$ for any fixed $\epsilon > 0$.

Given a binary tree, τ , with n leaves and height h , let $D_{opt}(\tau)$ denote the minimum communication delay needed to compute τ . It is shown that $\Omega(\log n) \leq D_{opt}(\tau) \leq O(\sqrt{n})$, and $\Omega(\sqrt{h}) \leq D_{opt} \leq O(h)$, all bounds being the best possible. We also present a simple polynomial algorithm that generates a schedule for computing τ with at most $2D_{opt}(\tau)$ delay.

It is shown that the a communication delay-computation time tradeoff given by Papadimitriou and Ullman for a diamond dag can be achieved for essentially two values of the computation time. We also present DAGs that exhibit proper tradeoffs for a substantial range of time.

1. INTRODUCTION

It is becoming abundantly clear that much of the complexity in parallel computing is due to the difficulty in communication rather than the computation itself. This issue is likely to become more severe as the number of processors increase. Researchers have recognized this fact, and studied the communication complexity of special purpose interconnection networks [A80, DGS83, PS81, T84, Y79] and of VLSI chips [AUY83, JK84, T79]. However, except for the paper by Papadimitriou and Ullman [PU87] which explores the communication complexity in rather broad terms, very little is known regarding the communication complexity of PRAMs. The purpose of this paper is to study

several aspects of this issue, and explore a few techniques that are useful in a better understanding of the communication requirements of PRAMs.

To capture the communication delay in multiprocessor systems, we model both the computational problem to be solved as well as the multiprocessing machine that is solving it. The model is similar to that used by Papadimitriou and Ullman [PU87] but since it differs in some essential features, it is described below.

The multiprocessing machine is taken to be a concurrent-read, exclusive write (CREW) PRAM in which each processor is provided with an unlimited amount of local random access memory. Processors can read simultaneously from the same location in the *global* memory, but two or more are not allowed to write simultaneously into the same location. The input variables are initially stored in the global memory, and the final outputs must be eventually stored in the global memory. The processors are MIMD, but work synchronously. In order to model the communication delay and computation time, it is convenient to restrict the machine such that at every time step, the processors do one of the following:

- (i) In one communication step, a processor can write and then read a word from the global memory, or it may do nothing.
- (ii) In a computation step, a processor can perform an operation (which we describe below) on at most two words that are present in its local memory.

The computation problem to be solved is usually presented as a directed acyclic graph (DAG) with its nodes corresponding to operations (on at most two variables) and its arcs corresponding to the values computed by performing such operations. Thus, the in-degree of each node is bounded by two, and if the out-degree of any node is more than one, then all its outgoing edges contain the same associated value. The nodes of the DAG with in-degree zero (i. e., the leaves) are available in the global memory at time $t = 0$, and likewise, the nodes of the DAG with out-degree zero (i. e., the roots), have to be eventually written into the global memory. Any processor correctly *computes* the given node if and only if the following conditions are satisfied:

- (a) The node is computed only after all its children have been computed.
- (b) A processor can perform the operation given at the node only if the corresponding values of its children reside in its local memory; it performs the operation on these values in one computation step. A processor can write a computed value into the global memory in a communication step, and can then read a value from the global memory in the same step.

For the computation of a given DAG, each of its nodes is assigned to one or more processors that *compute* that node. Hence, the problem of computing the DAG reduces


to that of finding a schedule for performing the operations given at the nodes which ensures that each node of the DAG is computed correctly. Figure 1 demonstrates a nine node DAG and a schedule that computes it in 5 communication steps and 3 computation steps. This DAG can also be computed in 4 (resp. 3) communication steps and 4 (resp. 7) computation steps.

Remark: Since many algorithms for computing a given problem can be converted into a DAG, one can usually obtain a plethora of DAGs for the same problem. Usually, we will be concerned with a fixed DAG for a given problem, although we will sometimes deviate and give lower bounds for a class of DAGs that can be used to solve problems such as matrix multiplication, or sorting.

Three kinds of resources are relevant -- the number of processors used in the multiprocessing machine, the overall computation time, and the overall communication delay. The computation time (communication delay, resp.) of a schedule, S , is the number of computation (communication, resp.) steps used by this schedule. The computation time for the DAG is the minimum computation time over all schedules that correctly compute the given DAG, and likewise, for communication delay. Note that the minimum communication delay and the minimum computation time for a given DAG may not be achievable by the same schedule.

2. MATRIX MULTIPLICATION AND RELATED RESULTS

Consider the computations of DAGs that are obtained when two $n \times n$ matrices are multiplied using scalar multiplications and additions only. We establish an optimal bound of $\Theta(n^2/p^{2/3})$ on communication delay where p denotes the number of processors and $p \leq n^3 / \log^{3/2} n$. Indeed, the communication delay of $\Theta(n^2/p^{2/3})$ may be contrasted with the optimal speed up of $\Theta(n^3/p)$ with respect to the number of computation steps for the same problem. The lower bound for matrix multiplication given in Theorem 2.3 can be extended to Boolean matrix multiplication (over a semi-ring containing AND and OR operations), to all pair shortest path problem (where only the semi-ring operations of MIN and + are allowed), and to transitive closure of matrices (if the closure is computed over a closed semi-ring with "+" and "x" as two operations). Furthermore, the upper bound of Theorem 2.1 can be generalized to demonstrate the optimality of the two bounds for Boolean Matrix Multiplication (within a constant factor) and to those for computing the all-pair shortest paths and transitive closure within $\log n$ factors. For $p \leq n^2 / \log n$, a PRAM with p processors can multiply an $n \times n$ matrix with an n -dimensional vector in $O(n^2/p)$ computation time and with $O(n^2/p)$ communication delay. Since all entries of matrix are initially stored in



Schedule: Comm. Step 1: P_1, P_2, P_3, P_4 read a .
 Comm. Step 2: P_1, P_2, P_3, P_4 read b .
 Comp. Step 1: P_1, P_2, P_3, P_4 compute c, d, e, f respectively.
 Comm. Step 3: P_2, P_4 write d, f respectively. P_1, P_3 read d, f respectively.
 Comp. Step 2: P_1 computes g , P_3 computes h .
 Comm. Step 4: P_3 writes h , P_1 reads it.
 Comp. Step 3: P_1 computes v .
 Comm. Step 5: P_1 writes v .

Figure 1: Example of a DAG and a schedule that computes it

the global memory, every PRAM algorithm would require $\Omega(n^2/p)$ communication delay. Furthermore, since a total of $\Omega(n^2)$ multiplications are required by every sequential program for matrix-vector multiplication (see [AHU74], pp. 428-435, for details), every PRAM algorithm would require $\Omega(n^2/p)$ computation steps and this establishes optimal bounds for matrix-vector multiplication.

Theorem 2.1: Two $n \times n$ matrices can be multiplied by p processors in $O((n^3/p) + \log n)$ computation time and $O((n^2/p^{2/3}) + \log n)$ communication delay using only $(\times, +)$.

Proof: Partition the two $n \times n$ matrices, A and B , in a natural manner, into $p^{2/3}$ submatrices of sizes $n/p^{1/3} \times n/p^{1/3}$ each, and for $1 \leq i, j \leq p^{1/3}$ let these matrices be denoted by $A_{i,j}$ and $B_{i,j}$ respectively. For $1 \leq i, j, k \leq p^{1/3}$ let the p processors read $A_{i,j}$ and $B_{j,k}$ such that each processor gets a unique pair of $(A_{i,j}, B_{j,k})$. Note that this global memory access takes $O(n^2/p^{2/3})$ communication delay. Now, compute the matrix $C_{i,j,k} = A_{i,j} \times B_{j,k}$ in $O(n^3/p)$ computation time. If $C = A \times B$ and if C is also partitioned, in a natural manner, into $p^{2/3}$ submatrices, $C_{i,k}$, for $1 \leq i, k \leq p^{1/3}$, then it is easy to see that $C_{i,k} = \sum C_{i,j,k}$. Now, summing these submatrices naively would require $O(n^2 \log n/p^{2/3})$ communication delay. However, using a *pipelining strategy*, which we give below, these submatrices can be added in $O((n^2/p^{2/3}) + \log n)$ communication delay. This yields an overall communication delay of $O((n^2/p^{2/3}) + \log n)$ and an overall computation time of $O((n^3/p) + \log n)$. For $p \leq n^3/\log^{3/2} n$, these bounds are $O(n^2/p^{2/3})$ and $O(n^3/p)$, respectively.

To sum the matrices $C_{i,j,k}$ for $1 \leq j \leq p^{1/3}$, and fixed i, k , use $2p^{1/3} - 1$ processors such that $p^{1/3}$ of these processors contain the submatrices $C_{i,j,k}$ for $1 \leq j \leq p^{1/3}$ and the $2p^{1/3} - 1$ processors compute a DAG in the form of a complete binary tree that has $p^{1/3}$ leaves and whose internal nodes are '+'. These processors can now add any $p^{1/3}$ elements present at the leaves in $O(\log p)$ time and, in fact, they can *pipeline* $n^2/p^{2/3}$ elements (that reside in the processors corresponding to each leaf) so that if at any instant, the processors that correspond to, say the l -th level of the tree are adding the elements which would eventually yield, say, the s -th entry of $C_{i,k}$ then at the same instant the processors which correspond to the $(l - 1)$ -th level, are adding the elements which would eventually yield the $(s + 1)$ -th entry of $C_{i,k}$. ■

For the lower bound on the communication delay for matrix multiplication, we need the following technical Lemma; the proof of this Lemma can be found in [HK81]:

Lemma 2.2: Let a PRAM compute the matrix multiplication of two $n \times n$ matrices, A, B using scalar additions and multiplications only. During this computation, if any processor reads at most s elements of A and B , and computes at most s partial sums of the product C , then this processor can compute a total of at most $2s^{3/2}$ multiplicative terms for these partial sums. ■

Theorem 2.3: Let G be any DAG with in-degree two that corresponds to an algorithm for multiplying two $n \times n$ matrices using $(+, \times)$ only and p processors. Then, any