

**A NOTE ON THE LOAD  
BALANCING PROBLEM FOR  
COARSE GRAINED HYPERCUBE  
DICTIONARY MACHINES**

Frank Dehne and Michel Gastaldo

SCS-TR-172, MAY 1990

School of Computer Science, Carleton University  
Ottawa, Canada, K1S 5B6

# A NOTE ON THE LOAD BALANCING PROBLEM FOR COARSE GRAINED HYPERCUBE DICTIONARY MACHINES \*

FRANK DEHNE<sup>1</sup> AND MICHEL GASTALDO<sup>2</sup>

<sup>1</sup> *Center for Parallel and Distributed Computing  
School of Computer Science, Carleton University, Ottawa, Canada K1S 5B6*

<sup>2</sup> *Laboratoire de l'Informatique du Parallelisme - IMAG  
Ecole Normale Supérieure de Lyon, 69364 Lyon cedex 07, France*

**Abstract.** The main problem for the design of dictionary machines on coarse grained hypercube multiprocessors, in comparison to the widely studied dictionary problem for fine grained hypercube multiprocessors, is that due to unequal distribution of the inserted and deleted records, the sizes of the sets stored at the individual processors may vary considerably. This problem, which is usually referred to as the *load balancing problem*, may lead to considerable degradation of the dictionary machine's performance. In this note we show that the load balancing problem for coarse grained hypercube dictionary machines can be solved with provable bounds on the sizes of the data sets, and with only little computational overhead.

## 1 INTRODUCTION

We present a solution to the load balancing problem arising in the design of dictionary machines for coarse grained hypercube multiprocessors.

The use of parallel architectures for implementing dictionary machines has been widely studied in the literature. Some authors design specialized architectures (e.g., [AK, BK, CI, F, ORS, SA]) while others use existing and widely available processor networks (e.g., [DS1, DS2, OB, SL]). Particular attention has been given to the problem of designing dictionary machines for hypercube multiprocessors [DS2, OB, SL]. A common feature of the presented solutions is that the records currently stored in the dictionary are stored in sorted order according to a pre-defined sequence of processors (called, e.g., "snake" and "sorted chain" in [DS2] and [OB], respectively). This sorted order is used to detect and avoid redundant deletions and duplicate insertions, respectively. The presented methods assume fine grained hypercube multiprocessors; i.e., every processor stores at most one record. In practice, one would often like to use coarse grained hypercube multiprocessors; that is, one would like to store a set of records at each processor. For the above methods, the snake (or sorted chain) of records would be replaced by a sorted sequence of sets of records (each of them usually represented as a balanced search tree). The main additional problem arising here, which is commonly

---

\* The first author's research is partially supported by the Natural Sciences and Engineering Research Council of Canada (Grant A9173).

The second author's research is partially supported by the Direction de la Recherche et des Etudes Techniques and the Programme de Recherche Coordonnées C3 (France).

referred to as the *load balancing problem*, is that due to unequal distribution of the inserted and deleted records, the sizes of the sets may vary considerably. This may lead to considerable degradation of the dictionary machine's performance.

In [OB] it is noted that one way of solving the problem is a periodic "shut down" of the dictionary machine followed by a global rebalancing phase (which is obviously not very desirable) and that "local balancing on the fly is considerable more difficult". They sketch a heuristic for shifting records between processors with unbalanced data sets, but do not give any performance guaranties.

In this note we show that the load balancing problem for coarse grained hypercube dictionary machines can be solved with provable bounds on the sizes of the data sets and with only little computational overhead. We present our result for the dictionary machine described in [DS2]; it can however be easily adapted to the solutions shown in [OB] and [SL].

## 2 BALANCING SEQUENCES OF 2-3 TREES FOR COARSE GRAINED HYPERCUBE DICTIONARY MACHINES

The hypercube dictionary machine presented in [DS2] consists of two structures, a snake and a broadcast net which are both embedded into a  $d$ -dimensional hypercube (with  $p=2^d$  processors) and operate on it simultaneously. The operations *insert*, *delete*, *find-min*, and *extract-min* are handled by the snake, while *search* operations are handled by the broadcast net (see [DS1, DS2] for more details). The snake, which is basically a systolic priority queue [L], is embedded into the hypercube as shown in Figure 1a and 1b (note that, in order to detect/avoid redundant deletions/duplicate insertions, the data elements are stored in sorted order with respect to the snake).




Figure 1a. Local Structure of the Snake in Each 4-dimensional Subcube.




Figure 1b. Global Structure of the Snake in a 7-dimensional Hypercube.

Let us now consider the coarse grained case,  $n \gg O(p)$ , where  $n$  is the number of data items stored in the dictionary. For the remainder we will assume that  $n \geq 2 p d$ . As indicated above, each processor will now store a set of data items as a balanced search tree; here, we will use 2-3 trees. Let us assume that at the beginning the global data structure is balanced, i.e. each processor stores the same number of entries. Note that these sets are sorted with respect to the snake ordering (i.e. all items in processor  $P_i$  are smaller than all items in  $P_j$ , if  $P_i$  is a predecessor of  $P_j$ ).



Figure 2. A Snake of Processors Each Storing a 2-3 Tree.

The load balancing problem consists now of keeping the sizes of the different search trees balanced while insertions and deletions into/from the data sets occur. In addition to having to shift around data items, a major problem is for the processors to determine whether they are unbalanced and to compute their correct size. Note that, since we allow duplicate insertions and redundant deletions, and because of the communication delay, it is impossible for each processor even to know, at any point in time, the total number of records currently stored in the dictionary.

Our approach to solving this problem is to compute at every processor approximate values of the current total number of records as well as the number of records to be shifted to the neighbors. We will ensure however that this information is precise enough to guarantee that the sizes of two arbitrary trees differ by at most  $O(\log p + \log m)$ , where  $m$  is the size of the largest tree in the dictionary machine. This, in turn, guarantees that the depths of two arbitrary trees differ by at most  $O(1)$ .

The general structure of our approach is as follows:

Assuming that the sizes of two trees differ by at most  $O(\log p + \log m)$ , we present an  $O(\log p + \log m)$  time hypercube algorithm to rebalance the trees to equal size (if the trees were not changed during this period). If this algorithm is concurrently executed with the general dictionary machine algorithm, then after  $O(\log p + \log m)$  steps the dictionary machine algorithm can only create  $O(\log p + \log m)$

imbalance between two arbitrary trees. Thus, if our rebalancing algorithm is always restarted immediately after its termination, the sizes of two trees will never differ by more than  $O(\log p + \log m)$ .

In the remainder, we will outline our  $O(\log p + \log m)$  time rebalancing algorithm. It consists of two phases: *approximate counting* and *shifting subtrees*. In our model, one exchange of an  $O(1)$  length message between two neighboring processors takes time  $O(1)$ , and we account for communication as well as for local processing time.

#### APPROXIMATE COUNTING

We assume that during the entire dictionary machine operation, every processor keeps a counter of the current size of its local tree. At the beginning of the first phase, every processor creates a copy ( $\sigma$ ) of that counter, and then for every processor  $p$  the following three numbers are computed:

- $n$  : the sum of the sizes ( $\sigma$ ) of all trees
- $n_p$  : the sum of the sizes ( $\sigma$ ) of all preceding trees (with respect to the snake), and
- $n_s$  : the sum of the sizes ( $\sigma$ ) of all succeeding trees (with respect to the snake).

We observe that, except for the 4-dimensional subcubes, the processor numbers of the processors in the snake sequence increase monotonically (see Figure 1b). Therefore, the above operation is essentially equivalent to a partial sum operation (see [KRS] for more details) and can hence be executed in time  $O(\log p)$ .

#### SHIFTING SUBTREES

With the above information, every processor determines the number  $s \leq O(\log p + \log m)$  of records it has to shift to the left or right, and from our initial assumption it follows that its tree contains at least  $s$  records. What remains to be shown is how to shift these records in time  $O(\log p + \log m)$  while maintaining the 2-3 tree structures at all processors.

Clearly, if we would shift every record individually and insert it into the neighboring tree, the required time would be  $O((\log p + \log m) \log m)$ ; i.e., it would be a factor of  $O(\log m)$  too long. We will therefore proceed as follows:

Assume that a processor  $P_i$  has to send  $s \leq O(\log p + \log m)$  records to its left neighbor and receive  $r \leq O(\log p + \log m)$  records from its right neighbor (in the snake):

- (1) Processor  $P_i$  cuts the  $s$  smallest records from its 2-3 tree  $T$ , rebuilds these  $s$  records into a 2-3 tree  $T'$ , and rebuilds the remainder into a valid 2-3 tree  $T''$ . It then sends the 2-3 tree  $T'$  to its left neighbor.
- (2) Processor  $P_i$  receives a 2-3 tree  $T'$  of size  $r$  from its right neighbor and merges it with its tree  $T''$  into a new 2-3 tree.

An example for the cutting and rebuilding operation referred to in Step 1 is given in Figure 3. We assume that at every node of the 2-3 tree we store, in addition to the minimum values of the middle and right subtree, also the minimum of the left subtree, the maximum of the right subtree, as well as the total number of leaves in all (at most three) subtrees. First, we navigate from the root down the tree to the  $s+1^{\text{st}}$  leaf (leaf number 10 in Figure 3). Then we rebuild both trees,  $T'$  and  $T''$ , in a bottom-up fashion similar to the standard insertion and deletion procedure for 2-3 trees. Note that at each level, only a constant number of links and nodes need to be shifted across the border of the two trees (or copies of nodes need to be created).

Summarizing we obtain that, assuming that the sizes of the subtrees differ by at most  $O(\log p + \log m)$ , the sequence of subtrees can be rebalanced in time  $O(\log p + \log m)$ . Note that, since the cutting and rebuilding operations (without transmission of the subtrees) can be executed in the same time as insertions or deletions into/from the 2-3 trees, problems with concurrent insertions and deletions are avoided.




Figure 3. An Example For Cutting and Rebuilding 2-3 Trees

## CONCLUSION

In this note we have shown that the load balancing problem for coarse grained hypercube dictionary machines can be solved with provable bounds on the sizes of the data sets and with only little computational overhead. Assuming that the sizes of two data sets differ by at most  $O(\log p + \log m)$ , we presented an  $O(\log p + \log m)$  time hypercube algorithm for rebalancing the sets to equal size (if they were not changed during this period). If this algorithm is concurrently executed with the general dictionary machine algorithm, then after  $O(\log p + \log m)$  steps the dictionary machine algorithm can only create  $O(\log p + \log m)$  imbalance between two arbitrary data sets. Thus, if our rebalancing algorithm is always restarted immediately after its termination, the sizes of two data sets will never differ by more than  $O(\log p + \log m)$ .

## REFERENCES

- [AK] M.J. Atallah and S.R. Kosaraju, "A generalized dictionary machine for VLSI", *IEEE Trans. on Computers* 34:2, 1985, pp. 151-155.
- [BK] J.L. Bentley, and H.T. Kung, "A tree machine for searching problems", Proc. 1979 Int. Conf. on Parallel Processing, 1979, May 1981, pp. 257-266.
- [CI] J.H. Chang, O.H. Ibarra, M.J. Chung, and K.K. Rao, "Systolic tree implementation of data structures", *IEEE Trans. on Computers* 37:6, 1988, pp. 727-735.
- [DS1] F. Dehne, N. Santoro, "Optimal VLSI dictionary machines on meshes", Proc. Int. Conf. on Parallel Processing, 1987, pp. 832-840.
- [DS2] F. Dehne, N. Santoro, "An optimal VLSI dictionary machine for hypercube architectures", Parallel and Distributed Algorithms, M. Cosnard (ed.), North Holland, 1989, pp. 137-144.
- [F] A.L. Fisher, "Dictionary machines with small number of processors", Proc. Int. Symp. on Computer Architecture, June 1984, pp. 151-156.
- [KRS] C.P. Kruskal, L. Rudolph, M. Snir, "The power of parallel prefix", Proc. Int. Conf. on Parallel Processing, 1985, pp.180-184.
- [L] C.E. Leiserson, "Systolic priority queues", Report CMU-CS-79-115, Carnegie-Mellon University, April 1979.
- [NS] D. Nassimi, S. Sahni, "Broadcasting data in SIMD Computers", *IEEE Trans. on Computers* 30:2, 1981, pp 101-107.
- [OB] A.R. Omondi, J. D. Brock, "Implementing a dictionary on hypercube machines", Proc. Int. Conf. on Parallel Processing, 1987, pp.707-709.
- [ORS] T.A. Ottman, A.L. Rosenberg, and L.J. Stockmeyer, "A dictionary machine for VLSI", *IEEE Trans. on Computers* 31:9, 1982, pp. 892-897.
- [SA] A.K. Somani and V.K. Agarwal, "An efficient unsorted VLSI dictionary machine", *IEEE Trans. on Computers* 34:9, 1985, pp. 841-852.
- [SL] A.M. Schwartz and M.C. Loui, "Dictionary machines on cube-class networks", *IEEE Trans. on Computers* 36:1, 1987, pp. 100-105.

**School of Computer Science, Carleton University**  
**Bibliography of Technical Reports**

- SCS-TR-131     **An Efficient Window Based System Based on Constraints**  
 Danny Epstein and Wilf R. LaLonde, March 1988.
- 
- SCS-TR-132     **Building a Backtracking Facility in Smalltalk Without Kernel Support**  
 See Third International Conference on OOPSLA, San Diego, Calif., Sept. '88.  
 Wilf R. LaLonde and Mark Van Gulik, March 1988.
- 
- SCS-TR-133     **NARM: The Design of a Neural Robot Arm Controller**  
 Daryl H. Graf and Wilf R. LaLonde, April 1988.
- 
- SCS-TR-134     **Separating a Polyhedron by One Translation from a Set of Obstacles**  
 Otto Nurmi and Jörg-R. Sack, December 1987.
- 
- SCS-TR-135     **An Optimal VLSI Dictionary Machine for Hypercube Architectures**  
 Frank Dehne and Nicola Santoro, April 1988.
- 
- SCS-TR-136     **Optimal Visibility Algorithms for Binary Images on the Hypercube**  
 Frank Dehne, Quoc T. Pham and Ivan Stojmenovic, April 1988.
- 
- SCS-TR-137     **An Efficient Computational Geometry Method for Detecting Dotted Lines in Noisy Images**  
 F. Dehne and L. Ficocelli, May 1988.
- 
- SCS-TR-138     **On Generating Random Permutations with Arbitrary Distributions**  
 B. J. Oommen and D.T.H. Ng, June 1988.
- 
- SCS-TR-139     **The Theory and Application of Uni-Dimensional Random Races With Probabilistic Handicaps**  
 D.T.H. Ng, B.J. Oommen and E.R. Hansen, June 1988.
- 
- SCS-TR-140     **Computing the Configuration Space of a Robot on a Mesh-of-Processors**  
 F. Dehne, A.-L. Hassenklover and J.-R. Sack, June 1988.
- 
- SCS-TR-141     **Graphically Defining Simulation Models of Concurrent Systems**  
 H. Glenn Brauen and John Neilson, September 1988
- 
- SCS-TR-142     **An Algorithm for Distributed Mutual Exclusion on Arbitrary Networks**  
 H. Glenn Brauen and John E. Neilson, September 1988
- 
- SCS-TR-143 to 146 are unavailable.
- 
- SCS-TR-147     **On Transparently Modifying Users' Query Distributions**  
 B.J. Oommen and D.T.H. Ng, November 1988
- 
- SCS-TR-148     **An  $O(N \log N)$  Algorithm for Computing a Link Center in a Simple Polygon**  
 H.N. Djidjev, A. Lingas and J.-R. Sack, July 1988  
 Available in STACS 89, 6th Annual Symposium on Theoretical Aspects of Computer Science,  
 Paderborn, FRG, February 16-18, 1989, Lecture Notes in Computer Science, Springer-Verlag No.  
 349
- 
- SCS-TR-149     **Smallscript: A User Programmable Framework Based on Smalltalk and Postscript**  
 Kevin Haaland and Dave Thomas, November 1988
- 
- SCS-TR-150     **A General Design Methodology for Dictionary Machines**  
 Frank Dehne and Nicola Santoro, February 1989
- 
- SCS-TR-151     **On Doubly Linked List ReOrganizing Heuristics**  
 D.T.H. Ng and B. John Oommen, February 1989
- 
- SCS-TR-152     **Implementing Data Structures on a Hypercube Multiprocessor, and Applications in Parallel Computational Geometry**  
 Frank Dehne and Andrew Rau-Chaplin, March 1989
-



- SCS-TR-153     **The Use of Chi-Squared Statistics in Determining Dependence Trees**  
R.S. Valiveti and B.J. Oommen, March 1989
- 
- SCS-TR-154     **Ideal List Organization for Stationary Environments**  
B. John Oommen and David T.H. Ng, March 1989
- 
- SCS-TR-155     **Hot-Spot Contention in Binary Hypercube Networks**  
Sivarama P. Dandamudi and Derek L. Eager, April 89
- 
- SCS-TR-156     **Some Issues in Hierarchical Interconnection Network Design**  
Sivarama P. Dandamudi and Derek L. Eager, April 1989
- 
- SCS-TR-157     **Discretized Pursuit Linear Reward-Inaction Automata**  
B.J. Oommen and Joseph K. Lancot, April 1989
- 
- SCS-TR-158     **Parallel Fractional Cascading on a Hypercube Multiprocessor**  
(revised) Frank Dehne, Afonso Ferreira and Andrew Rau-Chaplin, May 1989 (Revised April 1990)
- 
- SCS-TR-159     **Epsilon-Optimal Stubborn Learning Mechanisms**  
J.P.R. Christensen and B.J. Oommen, June 1989
- 
- SCS-TR-160     **Disassembling Two-Dimensional Composite Parts Via Translations**  
Doron Nussbaum and Jörg-R. Sack, June 1989
- 
- SCS-TR-161     **Recognizing Sources of Random Strings**  
(revised) R.S. Valiveti and B.J. Oommen, January 1990  
Revised version of SCS-TR-161 "On the Data Analysis of Random Permutations and its Application to Source Recognition", published June 1989
- 
- SCS-TR-162     **An Adaptive Learning Solution to the Keyboard Optimization Problem**  
B.J. Oommen, R.S. Valiveti and J. Zgierski, October 1989
- 
- SCS-TR-163     **Finding a Central Link Segment of a Simple Polygon in  $O(N \log N)$  Time**  
L.G. Alexandrov, H.N. Djidjev, J.-R. Sack, October 1989
- 
- SCS-TR-164     **A Survey of Algorithms for Handling Permutation Groups**  
M.D. Atkinson, January 1990
- 
- SCS-TR-165     **Key Exchange Using Chebychev Polynomials**  
M.D. Atkinson and Vincenzo Acciari, January 1990
- 
- SCS-TR-166     **Efficient Concurrency Control Protocols for B-tree Indexes**  
Ekow J. Otoo, January 1990
- 
- SCS-TR-167     **A Hierarchical Stochastic Automaton Solution to the Object Partitioning Problem**  
B.J. Oommen, January 1990
- 
- SCS-TR-168     **Adaptive List Organizing for Non-stationary Query Distributions. Part I: The Move-to-Front Rule**  
R.S. Valiveti and B.J. Oommen, January 1990
- 
- SCS-TR-169     **Trade-Offs in Non-Reversing Diameter**  
Hans L. Bodlaender, Gerard Tel and Nicola Santoro, February 1990
- 
- SCS-TR-170     **A Massively Parallel Knowledge-Base Server using a Hypercube Multiprocessor**  
Frank Dehne, Afonso Ferreira and Andrew Rau-Chaplin, April 1990
- 
- SCS-TR-171     **Parallel Processing of Quad Trees on the Hypercube (and PRAM)**  
Frank Dehne, Afonso Ferreira and Andrew Rau-Chaplin, April 1990
- 
- SCS-TR-172     **A Note on the Load Balancing Problem for Coarse Grained Hypercube Dictionary Machines**  
Frank Dehne and Michel Gastaldo, May 1990
-