

PARALLEL BRANCH AND BOUND ON FINE-GRAINED HYPERCUBE MULTIPROCESSORS*

FRANK DEHNE¹, AFONSO G. FERREIRA², AND ANDREW RAU-CHAPLIN¹

¹ Center for Parallel and Distributed Computing, School of Computer Science, Carleton University, Ottawa, Canada K1S 5B6

² Laboratoire de l'Informatique du Parallelisme - IMAG, Ecole Normale Supérieure de Lyon, 69364 Lyon cedex 07, France

Abstract. In this paper, we study parallel branch and bound on fine grained hypercube multiprocessors. Each processor in a fine grained system has only a very small amount of memory available. Therefore, current parallel branch and bound methods for coarse grained systems (≤ 1000 nodes) can not be applied, since all these methods assume that every processor stores the path from the node it is currently processing back to the node where the process was created (the back-up path). Furthermore, the much larger number of processors available in a fine grained system makes it imperative that global information (e.g., the current best solution) is continuously available at every processor; otherwise the amount of unnecessary search would become intolerable. We describe an efficient branch-and-bound algorithm for fine grained hypercube multiprocessors. Our method uses a global scheme where all processors collectively store all back-up paths such that each processor needs to store only a constant amount of information. At each iteration of the algorithm, all current nodes may decide whether they need to create new children, be pruned, or remain unchanged. We describe an algorithm that, based on these decisions, updates the current back-up paths and distributes global information in $O(\log m)$ steps, where m is the current number of nodes. This method also includes dynamic allocation of search processes to processors and provides optimal load balancing. Even if very drastic changes in the set of current nodes occur, our load balancing mechanism does not suffer any slow down.

1 INTRODUCTION

Branch-and-bound (B&B) search methods such as A* search, Alpha-Beta search, depth-first and best-first search are used in areas such as VLSI design, theorem proving, linear programming, chess playing, and many other Artificial Intelligence and Operations Research applications ([Win], [Nil]). Since most of the addressed problems are NP-hard, and therefore the size of the search space grows exponentially, many researchers have studied the parallelization of B&B methods (e.g. [HD], [New], [MC],

[UYII]). So far, studies have been aimed at using coarse grained multiprocessors for parallel B&B implementations (multiprocessors with a relatively small number - less than 1,000 -, of relatively powerful processors, each having a considerable amount of memory). Particular attention has recently been given to coarse-grained hypercube multiprocessors such as the FPS hypercube, NCUBE, or Intel iPSC (e.g. [LW], [Qui], [AC], [AM], [Fe], [MTM], [PW], [SGB]).

Parallel B&B methods for coarse grained multiprocessors aim at splitting the tree searched by the B&B method into subtrees, and having subtrees searched in parallel by different processors. The essential differences among the proposed coarse grained algorithms are how they deal with the two major problems arising in such an approach:

- (1) Load balancing: When a subproblem (a subtree to be searched) is assigned to an individual processor, the size of the subtree is not known in advance. Therefore, the sizes of the problems assigned to individual processors may vary significantly and this unbalanced distribution of work load may result in performance degeneration.
- (2) Global information: For sequential B&B methods, the pruning of the search tree often depends on global information such as the best solution found so far. On a multiprocessor, the distribution of global information causes additional computational overhead or, if information is not completely distributed, the parallel method may search more nodes than the respective sequential algorithm (unnecessary search).

In this paper, we study parallel B&B on fine grained hypercube multiprocessors (hypercubes with a large number - more than 10,000 - of small processors). The Connection Machine is an example of an existing fine grained

* This research was conducted while the second author was visiting Carleton University. The first author's research is partially supported by the Natural Sciences and Engineering Research Council of Canada (Grant A9173). The second author is currently on leave from the University of Sao Paulo (Brazil), project BID/USP, and partially supported by CAPES/COFECUB (Grant 503/86-9).

system. In the context of parallel B&B, fine grained multiprocessors obviously have the advantage of a larger number of processors and increased parallelism. However, each processor in a fine grained system has only a small constant amount of memory available. For any node p of the search tree let the *back-up path* of p be the path from p to the root of the search tree. In a fine grained multiprocessor, it is impossible for one processor to store the back-up path even for the one single node that it is currently examining. Therefore, the above mentioned coarse grained parallel B&B methods can not be applied to a fine grained system because all these methods assume that every processor stores the back-up path for its currently processed node (at least back to the tree node where the process was created).

Furthermore, the much larger number of processors (and, therefore, concurrent search processes) makes it imperative, for fine grained B&B algorithms, that global information (e.g., the current best solution) is continuously available at every processor since otherwise the amount of unnecessary search would become intolerable.

In the remainder of this paper we describe an efficient B&B algorithm for fine grained hypercube multiprocessors that solves the problem of storing all current back-up paths. Our method uses a global scheme where all processors collectively store all back-up paths such that each processor needs to store only a constant amount of information. Instead of storing the individual back-up paths, we store a *current back-up tree* which is defined by the union of the back-up paths of all search tree nodes currently under examination. (This also ensures that paths in the search tree shared by several back-up paths are stored only once, thus providing optimal storage utilization.) At each iteration of the algorithm, all nodes of the current back-up tree may decide whether they need to create new children, be pruned, or remain unchanged. We describe an algorithm that, based on these decisions, updates the current back-up tree and distributes global information in $O(\log m)$ steps, where m is the size of the current back-up tree. This method also includes a dynamic allocation mechanism for allocating search processes to processors, and provides optimal load balancing. Even if very drastic changes in the current back-up tree occur, the performance of our load balancing mechanism does not deteriorate.

Note that, the $O(\log m)$ overhead mentioned above is measured by considering only constant length messages, exchanged between adjacent processors, as $O(1)$ operations. In this model, an $O(\log m)$ overhead is very small; in fact, one single "PREF" operation (parallel inter-processor read) in Connection Machine *LISP has a larger time complexity.

The remainder of this paper is organized as follows: Section 2 introduces some notation regarding B&B algorithms as well as some standard hypercube operations that will be utilized in our algorithm. In Section 3, we will then present our parallel B&B algorithm for fine-grained hypercube multiprocessors. Section 4 concludes the paper.

2 DEFINITIONS AND BASIC HYPERCUBE OPERATIONS

In this section, we will first define a generalized B&B procedure in terms of five rules. By providing different procedures to implement these rules, users will be able to implement a wide range of B&B algorithms. We then review some basic standard hypercube operations which will be used in the remainder of this paper.

2.1 BRANCH AND BOUND

Branch-and-bound is a general technique for exploring search spaces; A*, Alpha-Beta, Hill-Climbing and Best-First search are well known instances of B&B algorithms.

In this paper we consider a general B&B algorithm defined by five rules; any particular B&B algorithm can be realized by providing a suitable implementation for each rule (see [AC]).

- Cost rule: Given a leaf node in the search tree and its parent's cost, this rule defines the cost up to and including the leaf.
- Bounding rule: Given a node in the search tree, this rule returns 1 if the node is no longer feasible (i.e., should be deleted) else returns a 0.
- Selection rule: Given a node in the search tree, this rule returns an integer specifying the number of children this node should create in this iteration.
- Expansion rule: Given a node, an integer specifying the number of children to be created, and a pointer to

where they should be created, this rule creates the new children.

- Termination rule: Given a set of global information, this rule returns true if a satisfactory solution has been found or all possibilities have been explored.

2.2 BASIC HYPERCUBE OPERATIONS

The branch-and-bound algorithm described in the next section uses slightly generalized versions of five well-defined hypercube operations. In addition to the registers listed below, implementations of these operations may require a constant number of auxiliary registers. In the following, for every register A available at every processor, $A(i)$ refers to register A at processor $PE(i)$. We assume a hypercube consisting of $N = 2^q$ processors.

Psum(source(i),result(i)): Every processor $PE(i)$ has some value stored in the register $source(i)$. This operation computes $result(i) := source(0) + source(1) + \dots + source(i)$ for each $PE(i)$. This is the standard partial sum operation in hypercubes and can be implemented in $O(\log N)$ time.

IdentifyBlock(block(i),endOfBlock(i)): A *block* of processors is defined by consecutive $PE(i)$'s having the same value stored in the register $block(i)$. For each $PE(i)$, this operation will then assign to $endOfBlock(i)$ the largest j such that $block(j)=block(i)$. Its time complexity is also $O(\log N)$, since it can be implemented by a "Concentrate" operation ([NS]) followed by a *RouteAndCopy* operation defined below.

BlockPsum(source(i),result(i),block(i)): This operation performs a partial sum within blocks only. Every $PE(i)$ will store in $result(i)$ the partial sum of the value stored in the register $source(i)$ within the block it belongs to. The *BlockPsum* operation can be implemented by a *Psum* operation followed by a *RouteAndCopy* operation. Hence its time complexity is $O(\log N)$.

Route(Reg₁(i),Reg₂(i),Dest(i),Cond(i)): Every processor $PE(i)$ has 2 data registers $Reg_1(i)$, $Reg_2(i)$, a destination register $Dest(i)$, and a boolean condition register $Cond(i)$. It is assumed that the destinations $Dest(i)$ are monotonic; i.e., if $i < j$ then $Dest(i) < Dest(j)$. This operation routes, for every processor $PE(i)$ with $Cond(i) = true$, the contents of register $Reg_1(i)$ to processor register $Reg_2(i)$ of

processor $PE(Dest(i))$; it can be implemented with an $O(\log N)$ time complexity by using a *Concentrate* operation followed by a *Distribute* operation described in [NS].

RouteAndCopy(Reg₁(i),Reg₂(i),Dest(i),Cond(i)): Under the same assumptions as for the *Route* operation, this operation routes, for every processor $PE(i)$ with $Cond(i) = true$, a copy of registers $Reg_1(i)$ to registers $Reg_2(Dest(i - 1) + 1)$, ..., $Reg_2(Dest(i))$, each; it can be implemented with an $O(\log(N))$ time complexity by using a *Concentrate* followed by a *Generalize* operation described in [NS].

3 AN ALGORITHM FOR PARALLEL BRANCH AND BOUND WITH GLOBAL INFORMATION ON A FINE-GRAINED HYPERCUBE

A branch and bound algorithm searches in the space of all the feasible solutions for a given problem. These feasible solutions are usually seen as a search tree S over the solution space. In parallel B&B algorithms, the search in S for an optimal or satisfactory solution is performed concurrently at several nodes of S . These nodes are referred to as *active nodes*. The subtree T of S defined by the union of the back-up paths of all active nodes is the *current back-up tree* as introduced in Section 1 (an example of a current back-up tree is shown in Figure 1). Note that we allow any node of the current back-up tree to be an active node, which allows e.g. the implementation of best-first search. For the remainder, m will refer to the size of the current back-up tree.




Figure 1. A Current Back-Up Tree (Bold Nodes).

In this section, we will describe a parallel B&B algorithm for fine grained hypercube multiprocessors. Our method stores the current back-up tree such that each processor needs to store only a constant amount of information. At each iteration of the algorithm, all active nodes of the current back-up tree decide whether they need to create new

children, be pruned, or remain unchanged. Based on these decisions, the algorithm described below updates the current back-up tree and distributes global information in $O(\log m)$ steps. This method also includes a dynamic allocation of search processes to processors, such that the work load balancing problem is solved optimally.

3.1 ALGORITHM OVERVIEW

We assume that every processor $PE(i)$ has a constant size register $n(i)$ to store one node of the current back-up tree T . Register $n(i)$ will contain fields $n.data(i)$, $n.cost(i)$, $n.parent(i)$, $n.children(i)$, $n.firstChild(i)$, $n.level(i)$ and $n.newChildren(i)$. Each of these fields stores a constant amount of data associated with the node $n(i)$, its cost, parent, number of children, position of first child, level, and number of new children, respectively. See Figure 2 for a list of registers required at every processor.

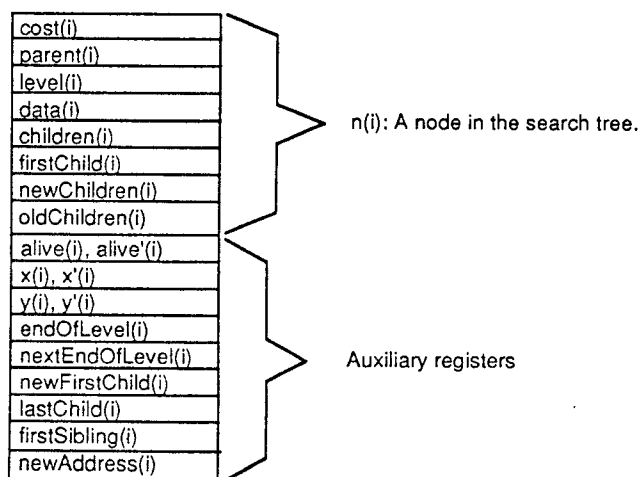


Figure 2. The Registers Required at Each Processor $PE(i)$.

The current back-up tree T is stored on the hypercube as follows:

Consider the level ordering of the nodes of T as shown in Figure 3. Each node v of T is stored in register $n(i)$ of processor $PE(i)$, where i is the index of v with respect to the level ordering of T .




Figure 3. Level Ordering of the Current Back-Up Tree

The global structure of the B&B algorithm is described in Figure 4. The back-up tree T starts as a single root node stored in register $n(0)$ of processor $PE(0)$. The main loop of the algorithm iterates until the termination rule returns true.

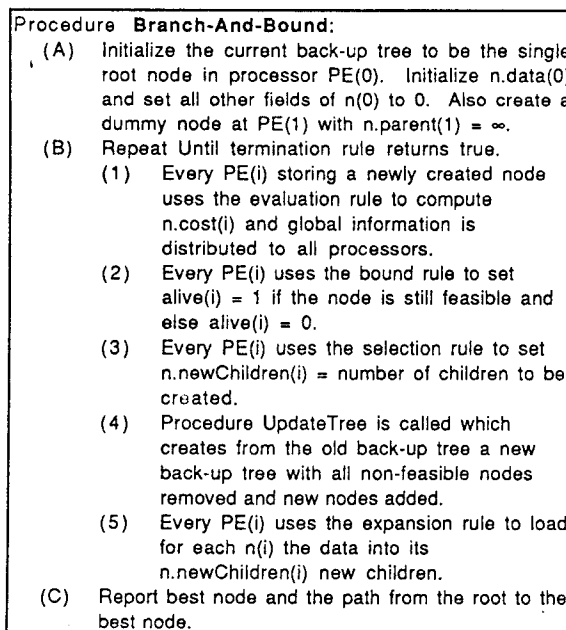


Figure 4. The Global Structure of the Branch-And-Bound Algorithm

In each iteration, the algorithm updates the tree by adding new nodes in a manner prescribed by the selection and expansion rules, while at the same time pruning the tree using the cost and bound rules. A single pass through the main loop consists of five steps. First, the evaluation rule is used to calculate the cost function for all the nodes that were created in the last iteration. The evaluation rule may also maintain global knowledge concerning the progress of the search. For example, after assigning new costs to nodes, the evaluation rule may calculate a global minimum cost to be

used by the bound and termination rules. The maintenance of such global knowledge does not increase the fundamental time complexity of the algorithm, since such global knowledge can be calculated by a global minimization and broadcast procedure in time $O(\log m)$, which is the same time complexity as the main body of the B&B algorithm.

The second step in the main loop of the algorithm uses the bound rule to identify nodes to be deleted in this iteration. Any particular implementation of the bound rule must never delete a parent node without also deleting its children. The parallel bounding of a large number of nodes is a particularly effective aspect of our fine-grained B&B algorithm, since with very low overhead a large number of nodes in the back-up tree may be deleted and the data structure compacted.

In the third step, the selection rule is used to assign to each node $n(i)$ an integer, $n.newChildren(i)$, that indicates the number of children the node needs to create during this iteration. The selection rule may also use global information to instruct each node exactly which of its possible children should be created.

Step 4, i.e. procedure `UpdateTree`, represents the main part of our algorithm. In Steps 1 through 3 all the information needed to extend the back-up tree to its next state has been collected. Now, the back-up tree must be transformed by deleting the bounded nodes and creating space for the new additional nodes. This operation is performed by procedure `UpdateTree` described in detail in the following Section 3.2.

Finally, in Step 5, the data for the new leaf nodes are created using the expansion rule. At this point each node has information about how many new children should be created, and about the address of the free processors allocated for them by procedure `UpdateTree`. Therefore, Step 5 can be implemented with an $O(\log m)$ time complexity by using a `RouteAndCopy` operation to copy the data of each node creating new children into the processors storing these children, and then creating for each child in parallel its actual data set locally at the respective processor.

3.2 UPDATING THE BACK-UP TREE

The core of our parallel B&B algorithm is procedure `UpdateTree` which updates the back-up tree after infeasible nodes have been deleted and new nodes have been added in the previous step. The main problem here is that for the new back-up tree, the nodes must again be stored by level number (to allow maximum storage space utilization and performance). This makes it necessary to compute for each node of the new back-up tree its new address and relocate the nodes to obtain the correct storage scheme. This relocation of tree nodes also provides an optimal task allocation mechanism for solving the load balancing problem.

The input to this procedure is stored in the two registers `alive(i)` and `n.newChildren(i)` at each processor. Register `alive(i)` is set to 1, if this node is not to be deleted, or 0 otherwise; `n.newChildren(i)` is set to the number of new children to be created for node $n(i)$.

The procedure `UpdateTree` calculates for each node where its children, if any, should be located on the hypercube after the update is completed. Each parent then broadcasts this information to all of its children. Since all nodes may move, all parent nodes must also broadcast their own new address to their children so that the children can update their parent pointers. When these steps have been completed, all non-deleted nodes are routed to their new locations. This route operation will leave the necessary spaces for the new children that are to be created by the expansion rule called in the main procedure. Procedure `UpdateTree` is described in detail in Figure 5.

```

Procedure UpdateTree:
1.0) BlockPsum(alive(i), alive'(i), n.parent(i))
    Route(alive'(i), n.children(i), n.parent(i),
          n.parent(i) ≠ n.parent(i + 1))
    Every PE(i): n.oldChildren(i) := n.children(i)
    Every PE(i): n.children(i) := n.children(i) +
          n.newChildren(i)
2.0) Psum(n.children(i), x')
    Each PE(i): x'(i) := x'(i) + 1
    IdentifyBlock(n.level(i), endOfLevel(i))
    Every PE(i): nextEndOfLevel(i) := endOfLevel(i+1)
    RouteAndCopy(x'(i), x(i), nextEndOfLevel(i),
                i = endOfLevel(i))
3.0) BlockPsum(n.children(i), y(i), level(i))
    Each PE(i): y(i) := y(i) - n.children(i)
4.0) Each PE(i): newFirstChild(i) := x(i) + y(i)
    Each PE(i): lastChild(i) := n.firstChild(i) +
          n.oldChildren(i) - 1
    RouteAndCopy(newFirstChild(i), firstSibling(i),
                lastChild(i), n.oldChildren(i) > 0)
    Every PE(i): newAddress(i) := firstSibling(i) +
          alive'(i) - 1
5.0) RouteAndCopy(newAddress(i), n.parent(i), lastChild(i),
                n.oldChildren(i) > 0)
    Each PE(i): n.firstChild := newFirstChild(i)
6.0) Route(n(i), n(i), newAddress(i), alive(i) = 1)

```

Figure 5. Procedure UpdateTree.

In order to obtain an illustration of the idea behind procedure UpdateTree, consider the node p in Figure 6. The address (processor number) of the first child of p, after the update is completed, is the sum of two numbers x and y, defined as follows: x is the number of nodes in the updated tree up to (and including) p's level, while y is the number of children of nodes to the left of p in p's level. The sum of these two numbers indicates the position of p's first child in the new back-up tree.

Procedure UpdateTree is composed of 6 steps. In Step 1, the number of children at each node is recalculated to account for the nodes that have been marked for deletion and the new nodes that are to be added. This operation is performed by first performing a partial sum over the number of alive nodes within each block of nodes that share the same parent. The result of the partial sum is then sent to each parent by its last child. Step 1 is completed by setting register n.children(i) of each node to the sum of its undeleted children and its new children to be created.




Figure 6. Illustration of the Address Calculation for the New Back-Up Tree.

In Step 2, the x value for each node is computed. The x value for any node at level k is one plus the sum of the number of children of all nodes up to (and including) the last node in level k-1. The x value for all nodes can therefore be calculated as follows. First, a partial sum on the number of children, n.children(i), of every node is computed and stored in register x'(i); x'(i) is then incremented by one to account for the root node. Then, the last node in each block of nodes sharing the same parent sends its x' value to its parent node.

In Step 3, the value of y for each node is calculated. This is performed by computing a partial sum on register n.children(i) for all nodes in the same level. In order not to count the children of a node in its own y(i) count, n.children(i) must then be subtracted from the value provided by the partial sum operation.

In Step 4, the new address of the first child of each node is computed and stored in register newFirstChild(i), using the x(i) and y(i) values previously calculated. This value is then broadcast to all the children of each node; since the back-up tree is stored on the hypercube in level order, all children of a node can calculate their new addresses by adding to the new address of their first sibling the number of other siblings that are before them.

Now, every node has its new address. In Step 5, all nodes with children broadcast their new address to their children, so that the children can update their parent pointers. Step 6 completes the UpdateTree procedure by moving every node to its new address.

4 CONCLUSION

In this paper we have presented an efficient B&B algorithm for fine grained hypercube multiprocessors. Our method uses a global storage allocation scheme where all processors collectively store all back-up paths such that each processor needs to store only a constant amount of information. At each iteration of the algorithm, all nodes of the current back-up tree may decide whether they need to create new children, be pruned, or remain unchanged. We have described an algorithm that, based on these decisions, updates the current back-up tree and distributes global information in $O(\log m)$ steps, where m is the current number of nodes. This method also provides a dynamic allocation mechanism that obtains optimal load balancing. Another important property of our method is that even if very drastic changes in the current back-up tree occur, the performance of the load balancing mechanism remains constant.

The proposed method is currently being implemented on the Connection Machine, and empirical performance test results will be available in the near future.

REFERENCES

- [AC] S.Anderson and M.C.Chen, Parallel branch and bound algorithms on the hypercube, in Hypercube Multiprocessors 1987, M.T.Heath, ed., SIAM Press, Philadelphia, PA, pp 309-317
- [AM] T.S.Abdelrahman and T.N.Mudge, Parallel branch and bound algorithms on hypercube multiprocessors, in the Proceedings of the Third Conference on Hypercube Concurrent Computers and Applications, 1988, G.Fox, ed., ACM Press, pp 1492-1499
- [Fel] E.W.Felten, Best-first branch and bound on a hypercube, in the Proceedings of the Third Conference on Hypercube Concurrent Computers and Applications, 1988, G.Fox, ed., ACM Press, pp 1500-1504
- [HD] S.R.Huang and L.S.Davis, Parallel iterative A* search : an admissible distributed heuristic search algorithm, International Joint Conference on Artificial Intelligence 89, Preprint
- [LW] G.J.Li and B.W.Wah, Coping with anomalies in parallel branch and bound algorithms, IEEE Trans. on Comp., vol. c-35, no. 6, Jun 86, pp 568-573
- [MC] A.Marsland and M.Campbell, Parallel search of strongly ordered game trees, Computing Surveys, vol. 14, no. 4, Dec 82, pp 533-551

- [New] M.Newborn, Unsynchronized iteratively deepening parallel alpha-beta search, IEEE Trans. on PAMI, vol. 10, no. 5, Sep 88, pp 687-694
- [Nil] Nils J. Nilsson, "Principles of Artificial Intelligence", Tioga Publishing Co, 1980
- [NS] D. Nassimi, S. Sahni, "Data broadcasting in SIMD computers", IEEE Trans. on Computers 30:2, 1981, pp. 101-106.
- [PW] R.P.Pargas and D.E.Wooster, Branch and bound algorithms on a hypercube, in the Proceedings of the Third Conference on Hypercube Concurrent Computers and Applications, 1988, G.Fox, ed., ACM Press, pp 1514-1519
- [Qui] M.J.Quinn, Implementing best-first branch and bound algorithms on hypercubes multicomputers, in Hypercube Multiprocessors 1987, M.T.Heath, ed., SIAM Press, Philadelphia, PA, pp 318-326
- [SGB] K.Scwan, J.Gawkowski and B.Blake, Process and workload migration for a parallel branch and bound algorithm on a hypercube multicomputer, in the Proceedings of the Third Conference on Hypercube Concurrent Computers and Applications, 1988, G.Fox, ed., ACM Press, pp 1520-1530
- [TM] F.S.Tsung and M.H.Ma, A dynamic load balancer for a parallel branch and bound algorithm, in the Proceedings of the Third Conference on Hypercube Concurrent Computers and Applications, 1988, G.Fox, ed., ACM Press, pp 1505-1513
- [UYII] H.Usui, M.Yamashita, M.Imai, T.Ibaraki, Parallel searches of game trees, Systems and Computers in Japan, vol. 18, no. 8, 1987, pp 97-109
- [Win] Patrick H. Winston, "Artificial Intelligence", Addison-Wesley, July 1984
- [WLY] B.W.Wah, G.J.Li and C.F.Yu, Multiprocessing of combinatorial search problems, Computer, Jun 85, pp 93108