PARALLEL SPLIT-LEVEL RELAXATION[†]

ASHOK SAMAL[*] and TOM HENDERSON

Department of Computer Science
University of Utah, Salt Lake City, UT 84112, USA

The goal of high level vision is to identify a set of regions in a given image. This has been called by various names: the scene labeling problem[1], the consistent labeling problem[2], the constraint satisfaction problem[3], Waltz filtering[4], the satisfying assignment problem[5], etc.

There are several approaches to solve this problem, including backtracking, graph matching and relaxation. A new method called *split-level relaxation,* which is based on discrete relaxation was proposed in Ref. 6. It takes care of multiple semantic constraints by considering each of them independently. The problem is known to be NP-complete, so it takes a long time to solve. With the advent of multiprocessors, it is now imperative to see if the problem can be solved faster in the average case.

In this paper we give a framework for solving the scene analysis problem in a parallel processing environment, using split-level relaxation. Experiments done on a multiprocessor show that it is indeed advantageous to use multiprocessors to solve this problem.

Keywords: Scene labeling; Discrete relaxation; Multiple semantic constraints.

1. INTRODUCTION

In this paper we give a framework for solving the scene analysis problem in a parallel processing environment, using split-level relaxation. Initial experiments are done on a multiprocessor to show its suitability. The paper is organized as follows. In Sect. 2, the consistent labeling problem is formally defined. The traditional ways of solving the problem are also described. The scene analysis problem, as used in the field of computer vision is formulated as a CLP (consistent labeling problem) in Sect. 3. The *split-level relaxation* process is described in Sect. 4. In Sect. 5 the features of the BBN Butterfly multiprocessor, which was used for our results are discussed. There are several ways one may extract parallelism from the split-level relaxation algorithm. They are discussed in Sect. 6. The details of the implementation and the results are given in Sect. 7. Finally, the conclusion and future research directions are briefly discussed in Sect. 8.

2. THE CONSISTENT LABELING PROBLEM

Before we define the problem formally, a few terms are explained.
- A *unit,* u_i, is an item that has to be assigned a value or a meaning. For example, in scene labeling, it could be a region obtained after segmenting the image.
- A *label,* l_i, is the value or the meaning that is associated to a unit. In the same context it could, for example, be grassland, water, etc. Usually there is only a fixed set of labels (called the domain of the unit) that can be assigned to a unit.

- A *unit-label*, $L_i = (u_i, l_i)$, is a pair consisting of a unit and its associated label.
- A *labeling*, $\mathbf{L}$, is a set $\{L_1, L_2, \ldots, L_i\}$ of unit-labels. A labeling is *complete* if there is a label for each of the the units; otherwise it is called a *partial* labeling.
- A unit can be assigned any value from its domain. In general, however, there are restrictions on the labels a set of units can have simultaneously in order to be consistent. We call these the *constraints*. They can be n-ary in general, but we restrict ourselves here to binary constraints only. The formulation can be extended to the general case easily. Thus, a constraint set $\mathbf{R}$ is a set, $\{r_1, r_2, \ldots, r_k\}$, where each r_i is a pair of unit-labels. If $((u_i, l_i), (u_j, l_j)) \in \mathbf{R}$, it means the unit u_i cannot have the label l_i, when u_j has the label l_j.
- A labeling $\mathbf{L} = \{L_1, L_2, \ldots, L_k\}$ is *consistent* iff $(L_i, L_j) \notin \mathbf{R}$, $\forall\, i, j \leq k$.

Using the above definitions, the consistent labeling problem, can be formulated as follows. Given,

- A set of units, $\mathbf{U} = \{u_1, u_2, \ldots, u_n\}$. Without loss of generality we assume that all the u_i's are distinct.
- Each unit u_i has an associated domain D_i, which forms the set of possible labels for the unit. We define $\mathbf{D} = \{D_1, D_2, \ldots, D_n\}$. Often, however, all the units take the values from the same domain, in which case $D_1 = D_2 = \ldots = D_n = D$.
- A constraint relation $\mathbf{R}$ which defines the constraints between the units and the labels.

the goal of the consistent labeling problem is to find a *complete, consistent* labeling. Thus,

$$\mathbf{CLP} \Leftrightarrow (\mathbf{U}, \mathbf{D}, \mathbf{R}).$$

There are several variations to this formulation. For example, instead of obtaining a complete labeling one might want the *best* labeling according to certain criteria. Another variation is to ask for the *maximal* labeling if a complete labeling is not possible. Sometimes all the consistent labelings may be required. The nature of the problem, however, remains unchanged.

2.1. Complexity of CLP

It has been proved that CLP is NP-complete[7]. Hence the algorithms that solve this problem are exponential in the worst case (unless P = NP). So, it is not important to compare the performance of the various algorithms in the worst case. What is more important is to compare their performance in the average case. Even here, it is difficult to analyze the algorithms. As Knuth[8] points out, it is difficult to analyze these algorithms without actually running them. Even then, the results may not be valid in a totally different problem domain. Gaschnig[5] took an experimental approach to study the performance of the algorithms. Similar results are also presented in Ref. 9.

2.2. Solutions to CLP

The consistent labeling problem can be solved in several ways. The simplest solution is the *generate-and-test* method. Here all possible labelings are enumerated and then the ones which are consistent are selected as the solutions. Clearly this method is very naive and inefficient when either $\mathbf{U}$ or D is a large set. For example, if $|\mathbf{U}| = 10$ and $|D| = 10$,

then the number of possible configurations is 10^{10}. In many cases the labels assigned to the first few units make the whole labeling inconsistent due to the nature of the constraints. If detected early, it can save a lot of computation.

The *standard backtracking* method takes advantage of the above observation. A single unit is assigned a label from its domain. Then another unit is chosen and it is given a label such that the new partial labeling is still consistent. If at some point it is not possible to find such a label, the process is backed up to the last unit that was assigned a label. It is given the next possible label and this process continues. If all the units are labeled without violating any constraint we have found a solution. If however we run out of labels at some point during this process, there is no solution.

Standard backtracking is more efficient than the generate-and-test method, but it is still not good enough for many practical applications. Several efforts have been made to improve the performance of the backgracking approach. Gaschnig[5] developed two new backtrack-type algorithms which performed better than the standard backtracking. In the *backmark* algorithm all redundant tests for checking the consistency between pairs of units are eliminated. The *backjump* algorithm is like the standard backtracking except it is possible to jump back more than one level when a failure occurs. Yet another approach called *forward checking* is proposed in Ref. 9. It performs better than the two previous algorithms in some cases, but the results are known to be valid only for the problems where each unit is constrained by all other units. Haralick et al.[10] have also described two *look-ahead* operators, Φ and Ψ, to reduce the computation during the backtracking process. Haralick and Shapiro[2,11] generalized these operators to incorporate arbitrary look-ahead.

Another approach was taken by Waltz[4]. The idea is to remove labels that are not consistent with any other label of other units. The removal of a label from one unit may in turn force other labels to be deleted. The changes are propagated and the process terminates when no further labels can be deleted, or a unit loses all its labels. In the latter case there is no solution to the problem. In the former case, when the process converges to a consistent set of labels, it is still necessary to go through a search procedure to get an unambiguous labeling. The rate of convergence depends on the nature of the constraints and the labels. A modified version of this method is described in Ref. 12. It is a parallel iterative procedure that allows probabilities to be associated with the labels.

Yet another approach to reduce the computation during the backtracking process was taken by Mackworth[3]. For binary constraints, the problem can be formulated using graphs, where nodes represent the units and the constraints are denoted by the arcs between the nodes. Each node of the graph has an associated label set. Such a network is called a network of constraints. Fundamental properties of such networks are explored in Ref. 7. Mackworth observed the thrashing behavior of the backtracking algorithms and developed three consistency tests: *node, arc,* and *path* consistency tests, which when enforced first may reduce the computation time of the backtrack programs drastically. Several algorithms to enforce this are also given in Ref. 3. Mohr and Henderson[13] developed an optimal algorithm for arc consistency and an improved algorithm for path consistency.

2.3. Applications of CLP

Many problems can be formulated as CLP. Gaschnig[5,14] used this approach to solve the N-queens problem, the Soma Cube, Instant Insanity, and the cryptarithmatic problem. Henderson and Davis[15,16] used it for syntactic shape analysis. Many graph problems, e.g. graph coloring, graph homomorphism, etc. can also be framed as CLP. As shown in the next section, the scene analysis problem can also be formulated as a CLP.

3. SCENE ANALYSIS AS A CLP

The goal of scene analysis is to locate and identify all the objects in a given image. There are several ways one might attempt to solve the problem. One approach is to first locate the features (e.g. holes, corners, etc.) in the image. Then these features and the constraints between them are used to identify the objects. Thus, the features in the scene constitute the set of units $\mathbf{U}$, the features of the possible objects in the scene are the labels D, and the relations between them define the constraint relation $\mathbf{R}$.

3.1. Models for the Objects

A model for an object is defined using the location and the orientation of the features, with respect to a fixed coordinate frame. These features define the label space or the domain of the units in the CLP. There are also constraints between these features, which in general may be n-ary. In practice, however, it is often sufficient to consider only the binary constraints. Thus, each model is completely defined by the features and the constraints, i.e. $M_i \Leftrightarrow (F_i^m, C_i^m)$, where F_i^m is the set of features and C_i^m is the set of constraints for the ith model, respectively.

It should be noted that the domain of a feature found in the image can contain a feature in any of the possible models. If the scene can have k possible objects, then the model set $\mathbf{M} = \{M_1, M_2, \ldots, M_k\}$. Another important observation is that even though all the constraints are binary, they are not always the same. For example, consider two holes H_1 and H_2 in an image, such that H_1 is both adjacent to and bigger than H_2. The two constraints can be expressed as binary predicates $adjacent(H_1, H_2)$, and $bigger(H_1, H_2)$. Here even though both are binary constraints, they are of different $type$. The $adjacent$ constraint is symmetric while the $bigger$ is not. The semantics of the two constraints are also very different. Hence, we group the constraints according to their types.

Since a constraint can be represented as a relation, the constraint set of a model can be expressed as a set of relations. Thus,

$$C_i^m = \{R_{i,1}^m, R_{i,2}^m, \ldots, R_{i,r}^m\},$$

where $R_{i,j}^m \subseteq F_i^m \times F_i^m$, $1 \leq j \leq r$ and r is the number of constraint types. The constraint set for the set of models is given by $C^m = \cup_{i=1}^k C_i^m$.

3.2. Information from the Image

Low level image processing operations are done in order to get all the features of the image. Now these features are used to do the recognition, i.e. to associate these feature

instances with a feature in any one of the models. The features obtained from the images are also constrained in a similar way to the features in the models, e.g. two line segments may be parallel, one may be longer than the other, etc. Here, also, one may find several kinds of constraints between the features, some of them may even be different from the constraints in the models. However, we choose only those kinds of constraints which are used to define the constraints for the models.

Hence, the information S, from the scene can be represented as:

$$S = (F^s, C^s),$$

where F^s is the set of features found in the image and $C^s = \{C_1^s, *C_2^s, \ldots, C_r^s\}$, $*C_i^s \subseteq F^s \times F^s$, $1 \leq i \leq r$. C^s denotes the set of constraints in the scene.

There is a direct correspondence between the two sets of constraints C^s and C^m. For each relation in $*C_i^s$ in C^s, there is an equivalent relation C_j^m in C^m which has the same physical interpretation, although the domains of the two relations are different.

3.3. Scene Analysis as CLP

In the previous sections we defined the information that can be obtained from the image and the models. Now using this we formulate the scene analysis problem as a CLP. The set of units that need to be labeled is the set of feature instances found in the image. The label set or the domain for each of the features may consist of any feature in any model with the same type. For example, if we find a hole in the image, then it could be any hole in any of the possible models. It is possible to narrow it down further by using their dimensions as well. See Ref. 6 for several ways to make it more efficient. For simplicity, however, we assume that the domain consists of the union of all the features of all the models. Also, we assume that no two features in the models have the same name. This can be easily avoided by using a naming scheme where the name of the model is also a part of the name of a feature. The constraints in this case are the constraints that are between the features in the models.

So the scene analysis (SA) problem can now be formulated as:

$$SA \Leftrightarrow (U, D, R),$$

where $U = F^s$, $D = \cup_{i=1}^k F_i^m$, and $R = (\cup_{i=1}^k C_i^m)$.

3.4. Example

We now present an example, to explain the above concepts. Figure 1 shows two simple industrial parts similar to the ones in Ref. 17. We will refer to them as $Part_1$ and $Part_2$. They are not as complicated as many other parts, but they are only used for illustration. We use only the boundary edges as the features. Other useful features can be holes, corners, etc.

Fig. 1. Two industrial parts used in the example.

$\mathbf{M} = \{Part_1, \ Part_2\}$

$Part_1 = (F_1^m, \ C_1^m)$ and $Part_2 = (F_2^m, \ C_2^m)$

$F_1^m = \{Y_1Y_2, \ Y_2Y_3, \ X_4Y_3, \ \ldots, \ Y_1Y_6, \ Y_1Y_2\}$

$C_1^m = \{parallel, \ perpendicular, \ longer, \ equal\text{-}length\}$

$F_2^m = \{Z_1Z_2, \ Z_2Z_3, \ Z_3Z_4, \ Z_4Z_1\}$

$C_1^m = \{parallel, \ perpendicular, \ longer, \ equal\text{-}length\}$

We have four relations to express the constraints in the model for $Part_1$. These four relations are given in Table 1. These relations are not complete, but give an idea of what they look like. $Part_2$ has similar relations to express its constraints.

Table 1. Constraint relations in $Part_1$.

Parallel	Perpendicular	Longer	Equal-length
$(X_3Y_4, \ X_1Y_8)$	$(Y_1Y_2, \ Y_2Y_3)$	$(Y_1Y_6, \ Y_1Y_2)$	$(Y_1Y_2, \ Y_5Y_6)$
$(X_1Y_8, \ X_2Y_7)$	$(Y_2Y_3, \ X_4Y_3)$	$(Y_1Y_6, \ X_3X_4)$	$(X_3X_4, \ X_1X_2)$
$(X_2Y_7, \ Y_5Y_6)$	$(X_4Y_3, \ X_3X_4)$	$(Y_1Y_2, \ Y_2Y_3)$	$(X_4Y_3, \ X_3Y_4)$
$(Y_1Y_6, \ X_3X_4)$	$(X_3X_4, \ X_3Y_4)$	$(X_3X_4, \ Y_2Y_3)$	$(X_1Y_8, \ X_2Y_7)$

4. SPLIT-LEVEL RELAXATION

One of our main motivations for modifying the structure of the standard discrete relaxation procedure is that there is no way of using positive information. In relaxation all the labels of a node are treated uniformly. For example, suppose there is evidence that some labels of a unit are more likely than the others. The relaxation process works just the same as the case where there is no such evidence. One could go to a stochastic model, but it unduly complicates the process. We will take another approach without assigning a

probability value to each label of each unit. We briefly describe the method here. For details see Refs. 6 and 18.

4.1. Network Model for Relaxation

Here the underlying network/graph model for split-level relaxation process is explained. As mentioned before, the units are represented by the nodes of the graph and the arcs denote the constraints. In that respect, our model is similar to that explained in Ref. 3. However, there are several major differences. What we have here is conceptually closer to a set of graphs rather than a single graph.

Each of the graphs models a single relation, which represents only one type of constraint. So, if there are r types of constraints between the features in the image, then there are r corresponding graphs. There are also r graphs for each of the models that is considered. We call these graphs image and model graphs, respectively.

4.1.1. Model graphs

Each model graph G_i^m is a composite graph consisting of a set of graphs each corresponding to a single constraint. Thus,

$$G_i^m = \{G_{i,1}^m, \ G_{i,2}^m, \ \ldots, \ G_{i,r}^m\},$$

where r is the number of constraint types in the graph. The nodes of the graph $G_{i,j}^m$ are the set of features in the model, and the arcs are represented by the constraints. So,

$$G_{i,j}^m = \ <F_i^m, \ E_{i,j}^m>,$$

where $(x, y) \in E_{i,j}^m$ iff $(x, y) \in C_{i,j}^m$. Recall that $C_{i,j}^m$ is the relation that represents the jth constraint in the ith model.

We also denote G^m to denote the composite graph consisting of all the model graphs in the graph. So,

$$G^m = \{G_1^m, \ G_2^m, \ \ldots, \ G_k^m\}.$$

4.1.2. Image graphs

In contrast to the model graphs, there is only one image graph, although it is also possible to have multiple image graphs if we are working on multiple images at the same time, or labeling different parts of one image simultaneously, in a disjoint fashion (see Ref. 6). The composite image graph G^s consists of a set of graphs corresponding to the constraint types. Thus,

$$G^s = \{G_1^s, \ G_2^s, \ \ldots, \ G_r^s\},$$

where r is the number of constraint types in the graph. Similarly,

$$G_i^s = \ <F^s, \ E_i^s>,$$

where $(x, \ y) \in E_i^s$ iff $(x, \ y) \in C_i^s$.

4.2. Split-Level Relaxation

Since our network model is different from the standard network models, the relaxation process is also very different. The first step in this process is to build all the graphs: model and image graphs. Also, the possible labels for the units in the image graphs (see Sect. 4.1.2.) are also associated with the corresponding units. It is important to note that although the arcs of the image graphs are different, their nodes are the same. Hence they can be shared in an actual implementation. The topology of all the graphs remains the same during the relaxation process. Only the label sets of the units in the image graphs may change. The model graphs remain totally unchanged; they are essentially used to check for the validity of constraints between the features in the models.

After the graphs are built, the next step is to enforce the node, arc, and path consistency checks in the image graphs. This is where the system lends itself to parallel execution which will be explored further in Sect. 6. Here only the structure of the split-level relaxation algorithm is described. We divide the nodes of the graph into two categories: *strong* and *weak* nodes. In contrast, standard relaxation treats all the nodes equally. The strong nodes signify positive information, and always remain strong during the whole process. The weak nodes, however, may be elevated to the strong status. During the relaxation process the strong nodes can affect the label sets of the other nodes, both strong and weak. The weak nodes, however, can affect only the label sets of weak nodes. This causes the bad information, which is in the form of weak nodes, to have much less effect than the strong nodes, which signify positive information. After some initial iterations, the weak nodes that survive are elevated to strong nodes. The ones that do not survive the initial iterations are not considered further and are treated as features that are not part of any model. The enforcement of consistency is done until there is no change in the label set of the nodes.

Other ways to solve the problem of "weak" nodes deleting the labels of the "strong" nodes are explained in Ref. 6. Several issues considering the efficiency of the algorithms are also discussed there.

5.   THE BBN BUTTERFLY PARALLEL PROCESSOR

The results presented in this paper are obtained using the BBN Butterfly parallel processor. A brief discussion of the main features of the machine is given in this section. The Butterfly is a shared-memory MIMD machine capable of having up to 256 processors. All the processors in the machine are identical. Each processor may run its own program independently of the other processors. The memory is shared between all the processors and may be used to communicate between the processors. The memory is physically partitioned among the processors, but each processor can access the memory of every other processor through a switch. Thus even though it is tightly coupled, it is not a true shared-memory multiprocessor. In contrast, the Cosmic Cube[19] is loosely coupled, where a processor can access only the memory that is local to it. An example of a truly shared-memory machine is the Sequent[20].

5.1. Hardware

The Butterfly hardware consists of two main subsystems—the processor nodes and the butterfly switch. The processor nodes are responsible for the computing job of the machine, while the switch forms the communication system of the machine, and is responsible for the communication between the processor nodes.

The Butterfly parallel processor can have up to 256 processor nodes. Each processor node consists of an MC68020 microprocessor, a processor node controller (PNC), and up to 4 megabytes of memory. The machine used to obtain the results in this paper has 18 processors, with all but two nodes with 1MB of memory. Each processor node has a M68881 floating point co-processor along with the MC68020 microprocessor.

The PNC occupies a central role in the machine architecture. All the accesses to the memory including the local memory accesses are directed to the PNC. It uses the switch for the remote memory references. In addition it performs several other operations, e.g. atomic operations, I/O interfaces, etc.

All the processor nodes are connected through a switching network called the Butterfly switch. It is a self-routing, packet switching Ω network. One important characteristic of the Butterfly switch is that the number of switching elements grows as $N \log N$, where N is the number of processing elements. The bandwidth of the network grows linearly with the number of processors. The maximum data transfer rate between two processor nodes is 32Mbit/second per channel. The ratio of time taken for a remote memory access to a local memory access is roughly 5 to 1 (see Ref. 21 for some tests and precise numbers).

5.2. Software

The application programs run on the Butterfly under the Chrysalis Operating System. Currently C and a few other languages are supported. All the low level Chrysalis subroutines can be called from the application program. These include primitives for process creation, synchronization, creation of memory objects, creation and maintainance of queues, etc.

An application library called the Uniform System is also provided. It supports a methodology for programming, where all the processors share a common address space. It is written on top of the Chrysalis system. The application is structured such that all the available processors are used. While it is convenient to use for many applications, it is not very flexible. The programs written for obtaining the results in this paper are written using the low level Chrysalis functions directly. The features that have direct bearing on our application are discussed in Sect. 7.

The memory management system of the Butterfly has important consequences in the design of algorithms. The first Butterfly was designed using the MC68000, which supports only 24 bit addresses. Now even though MC68020 (which supports 32 bit addresses) processors are used, only 24 bit addresses can be used because of other hardware restrictions. So the problem here is to map the 24 bit virtual address to a 32 bit physical address. The details of this scheme are discussed elsewhere[22,23,24].

6. PARALLELISM IN SPLIT-LEVEL RELAXATION

In this section we explain how the split-level relaxation works in a parallel processing framework. It is shown that it is really advantageous to run this on a parallel computer since there is a tremendous amount of parallelism inherent in the algorithm. There are actually three levels of concurrency that may be exploited during this process.

At the topmost level, it is possible to look for objects in different parts of the same image. It can be done in an efficient manner as explained in Ref. 6. We have not tried to exploit this level of parallelism. Split-level relaxation itself lends very well to the parallel execution at the top level, since each of the graphs can be worked on independently. Finally, there is lot of parallelism in the node, arc, and path consistency algorithms that can also be exploited effectively.

6.1.   Parallelism in the Relaxation Process

We essentially have a set of graphs which are independent, since they represent different types of constraints, although they share a set of common nodes. Thus, the consistency (node, arc and path) tests can be performed in each of the graphs simultaneously. In the example in Sect. 3.4 we have four types of constraints, and hence we can have processes working concurrently trying to enforce consistency in the four graphs: parallel, perpendicular, longer and equal-length. Figure 2 gives a top level description of this process.

The **split-level** process creates r-1 processes, where r is the number of constraint types. It synchronizes the children, and also does work on one graph. The procedure **consistency**(i) does the node, arc, and path consistency in the ith graph. Three shared variables are used for synchronization of the processes: *counter, state* and *status*. If any label is deleted while doing the consistency checks, **consistency** procedure sets the *status* flag to CHANGE. *AtomicAdd* does the addition atomically. This is a synchronous algorithm and it has to be synchronized after each iteration. It can be proved that there is no deadlock in this scheme, although there are some spin waits.

It is easy to show that this scheme is indeed correct, i.e. it always terminates and produces the correct result.

Theorem 1.   *If a label* l *is removed from the label set of any node* u *during the parallel split-level relaxation process, then the node* u *cannot have the label* l *in any complete consistent labeling.*

Proof.   If the label l is removed by some process i during the process, then l is not a consistent label in the ith graph. Since it is necessary for the label to be consistent in all the graphs (i.e. to obey all the constraints of all the types), it cannot be a part of complete consistent labeling. It does not make a difference when a label gets deleted during the relaxation nor by which process.

Theorem 2.   *If a label* l *is not removed from the label set of any node* u *during the parallel split-level relaxation process, then it does not violate any consistency tests.*

Proof.   It is obvious from the structure of the relaxation process.

```
split-level ()
{
   build-graphs();                       /* Build model and image graphs */
   state = INIT;                         /* Initialize state */
   status = CHANGE;                      /* Set the status value */
   for (i = 1; i < r; i + +)
      create-process(work-graph,i);      /* Create one process per graph */
   while (status ! = NOCHANGE) {         /* Repeat until no change */
      counter = 0;                       /* Initialize counter to zero */
      status = NOCHANGE;                 /* Reset the status */
      state = WORK;                      /* let the children work */
      while (counter < r - 1);           /* Wait for all children to start */
      state = SYNC;                      /* this is for synchronization */
      counter = 0;                       /* reinitialize counter */
      consistency(0);                    /* work on own graph */
      while (counter < r - 1);           /* Wait for children to sync */
   }
   state = QUIT;                         /* Kill child processes */
}


work-graph(graph-type)
{
   while (state = = INIT);               /* Wait for signal */
   while (state != QUIT) {               /* Work until parent signals QUIT */
      while (state != WORK);             /* Work for WORK signal */
      consistency(graph-type);           /* consistency checks */
      AtomicAdd(counter,1);              /* Update counter */
      while (state == WORK);             /* Wait for others to finish */
      AtomicAdd(counter,1);              /* used for synchronization */
   }
}
```

Fig. 2. Structure of parallel split-level relaxation.

The two theorems prove that the parallel split-level relaxation process will produce the exact same result as the sequential one. However, it will, in general, arrive at the result faster.

6.2. Parallelism in the Consistency Algorithms

In this section it is shown that there is a lot of parallelism in the consistency algorithms. It is straightforward to see that the node consistency for a network of n nodes and a labels can be done in $O(1)$ time, if na processors are available. This can be done by doing the consistency checks for each label in each node in parallel. There is also parallelism in the path consistency algorithms, but we have not tried to exploit it, since in general it is not clear if the cost of using it is worth the effort. Most of our effort has been devoted to finding and exploiting the parallelism in the arc consistency algorithms.

Parallelism in arc consistency algorithms has been explored in detail in Refs. 25 and 26. Here we present only the relevant results and a sample parallel algorithm for achieving arc consistency.

It has been proved in Ref. 26 that any parallel algorithm doing arc consistency (using polynomial number of processors) will at least take $O(na)$ time, where n and a are the number of units and the number of labels per unit, respectively. It is also shown that the parallel AC1, AC3 and AC4 (called PAC1, PAC3 and PAC4, respectively) are all $O(na)$ algorithms. However, their performance in the average case may be very different. Their computational requirements are also very different. The PAC1 (parallel algorithm for AC1) always needs $O(ea^2)$ (e is the number of arcs in the graph) processors to attain maximum speedup, but the PAC3 and PAC4 algorithm need fewer processors on the average. In the worst case, however, they all need the same number of processors.

Next we briefly sketch the structure of a parallel AC3. It is shown in Fig. 3. The difference between AC1 and AC3 is that only the nodes whose label sets are changed contribute to the arcs to be examined next. So, the number of arcs which could be viewed as a task, varies from one iteration to the next. Thus a queue which maintains the list of current tasks is very suitable. However, the *enqueue* and *dequeue* operations should be done atomically to maintain the correct semantics.

```
consistency (graph-type)
{
  build-graphs();
  state = WORK;                            /* Initialize state */
  TaskQ = Make-Queue();
  initialize-queue();                      /* add the initial task descriptions */
  for (i = ; i < n; i + +)
    create-process(revise,i,TaskQ);        /* Create one process per processor */
  while (state != QUIT) {                   /* Work until Queue is empty */
    task = Wait-Queue(TaskQ);              /* Wait for a task description */
    change = work(task);                   /* work on the task */
    if (change)
      add-tasks-to-queue(task,TaskQ);      /* Add new tasks to queue */
    if empty(TaskQ) state = QUIT;          /* No more tasks; so quit */
  }
}


revise(index, TaskQ)
{
    while (state != QUIT) {                 /* Wait until parent signals QUIT */
    task = Wait-Queue(TaskQ);              /* Wait for a task description */
    change = work(task);                   /* work on the task */
    if (change)
        add-tasks-to-queue (task, TaskQ);  /* Add new tasks to queue */
    }
}
```

Fig. 3. Structure of a parallel AC3 algorithm.

It should be noted that the termination criterion is slightly complicated. The processes should stop if and only if the queue is empty and no process is working (or will add to the queue). The *work* procedure essentially does the consistency checks as per the specification of the task. If there are any changes then *add-tasks-to-queue* adds new tasks to the queue. The parent after creating the child processes also waits for the tasks and

works on them as well. Several details which are necessary for the algorithm to work correctly on a multiprocessor are omitted, but this should give an idea of the overall structure of the algorithm.

7. IMPLEMENTATION DETAILS

In this section we describe the details of the implementation of the split-level relaxation on the BBN Butterfly. Our machine has 18 nodes of which two nodes are used by the system and hence are unavailable. All the code is written using the low level Chrysalis primitive functions. We chose to do that in order to have more control over the process creation, task partitioning, memory management, etc.

There are several ways to divide the tasks in order to exploit the parallelism in the system, since it is so rich in parallelism. However, since we have only 16 processors available, we take a multi-sequential programming approach. The basic idea here is that we create exactly one process on each processor which runs until the whole program is terminated. Since process creation is expensive on the Butterfly, it is a good design decision. Each process looks like the **revise** process in Fig. 3. It goes through several stages: work, synchronization, etc. One could design an asynchronous version of this algorithm, but our implementation is a synchronous one.

We use a *queue* (DualQ, in Butterfly terminology) to keep all the tasks that need to be executed. Each process waits for a task in the queue. Once it finds a task, it dequeues it, and does the work specified by the task. The granularity of the task should be carefully chosen. If the granularity is too large then the time taken for synchronization at the end of each iteration may be large. If on the other hand it is too fine, the overhead costs in waiting for the queue etc. may be overwhelming. There are several ways to divide the tasks in our algorithms. We could make each task do the consistency checks on one whole graph. This is too large grained. On the other hand if we go down to the level of doing consistency checks for each label on each node it becomes too fine grained. We take an intermediate approach. The duty of each task in our system is to enforce consistency along one arc in a particular type of graph. So, if there are r constraint-types and e arcs, there are re tasks in the system. Also, we chose to use parallel AC1 as the underlying consistency alogrithm. We did not attmept do to path consistency.

Several low level Chrysalis function calls are very useful. Among them are Atomic_Add to do the atomic operations, Make_DualQ, Enq_DualQ, and Wait_DualQ to create and manipulate queues, etc. The memory is partitioned over all the processors to reduce memory contention. We use only one centralized queue whose address is informed to all the processes. This will be a bottleneck for a large multiprocessor, but is not a big problem in our implementation.

7.1. Test Problems

The algorithms are tested on three types of graphs: cyclic graphs, complete graphs and random graphs. The cyclic graph problem is described in Ref. 25. The nodes of the graph are connected to form a cycle, see Fig. 4. The label set of node i is given by $\{x : jn + i, 1 \leq j \leq a\}$, where there are n nodes and a labels per node. If the constraint used is the

greater-than relation, only one label is removed from the network in one iteration and hence it takes *na* iterations to converge. There is no solution to this problem and hence, the label sets all go to nil. In addition to the *greater-than* relation we use two more constraint types. The result is still the same, but the convergence to the solution is faster.


Fig. 4. A network of constraints.

The complete graph is used to model the N-queens problem, using the row placement strategy. The possible labels for each queen are $1, 2, \ldots, n$. The constraint is that no two queens should attack each other. This can be broken up into two constraints as follows. If two queens in ith and jth row are in columns k and l, respectively, the two constraints are: $k \neq l$ and $|i - j| \neq |k - l|$. We use these two constraints for this problem. If we do not constrain the search space further, nothing will happen, since for every position of queen i there is at least one position for queen j which is compatible. We chose to fix the position of two queens (nth and $n - 1$st) such that they conflict and hence there is no solution. This forces the label sets of all units to go to nil.

The underlying graph in the N-queens problem is a complete graph; hence the number of arcs is $O(n^2)$. In the case of the cyclic graph the number of arcs is $O(n)$. So we decided to test some graphs whose number of arcs is between the two. We generate a random graph with number of arcs $= O(n^{1.5})$. They are connected depending on the results of a random event, in this case on the value of a uniform random number generated.

All these problems are tried for several problem sizes (number of units). The results are shown for $n = 20, 30, 40$ and 50. The number of labels in N-queens is automatically fixed by n. In the case of the other two problems we chose it to be the same as n.

7.2. Results

The major goal was to see how much benefit is obtained by running the split-level relaxation algorithm on the Butterfly. So the major yardstick we use here is *speedup*. It is defined as follows:

$$S_k(A) = \frac{Time \ taken \ by \ an \ algorithm \ A \ on \ 1 \ processor}{Time \ taken \ by \ the \ same \ algorithm \ A \ on \ k \ processors}.$$

In the ideal case one would like to get a linear speedup with a slope of one. It is also very hard to obtain in practice. The speedup as a function of the number of processors for the three problems is given in Figs. 5, 6 and 7. In the random graph problem, the maximum speedup obtained is about 13, using all 16 nodes. In the other two problems we achieve what is called *super-linear* speedup, i.e. the speedup is greater than the number of processors. We achieve this in the cyclic graph problem and the complete graph problem. It happens in these cases because of combinatorial implosion. Since the processes are working concurrently, a change effected by any one of them is immediately seen by the others. This may save some work for the parallel algorithm which the sequential algorithm has to do. In the case of complete graph, we achieve super-linear speedup only for $n = 50$ using 15 or 16 processors. See Ref. 27 for detailed discussion.

For smaller problem sizes, the speedup numbers are not very good since there are fewer tasks and the time spent in synchronizing is high, particularly as the number of processors goes up. That explains the drastic drop in performance of the algorithm for $n = 20$. As the problem size increases however, the performance improves considerably.


Fig. 5. Speedup vs. number of processors for the cyclic graph.

Fig. 6. Speedup vs. number of processors for the complete graph.

Fig. 7. Speedup vs. number of processors for the random graph.

8. CONCLUSION AND FUTURE RESEARCH

In this paper we analyzed the split-level relaxation technique in a parallel processing framework. It is shown that there is much parallelism inherent in the algorithm that can be exploited. The algorithms are implemented on an actual multiprocessor and the results confirm this.

Although the results look very encouraging, the implementation can be made more efficient. There are several ways to improve upon it. Using multiple queues instead of one centralized queue will reduce the memory contention, particularly in large multi-processors. Also, switching to an asynchronous model should provide a major improvement in the performance. Another direction for future research is to explore parallelism in the path consistency algorithms.

REFERENCES

1. Dana H. Ballard and Cristopher M. Brown, *Computer Vision*, Prentice Hall, Inc., 1982.
2. Robert M. Haralick and Linda G. Shapiro, "The consistent labeling problem: Part I", *IEEE Transactions On Pattern Analysis And Machine Intelligence* **1** (1979) 173–184.
3. Alan K. Mackworth, "Consistency in network of relations", *Artificial Intelligence* **8** (1977) 99–118.
4. David Waltz, "Understanding line drawings of scences with shadows" in *The Psychology of Computer Vision* McGraw-Hill, 1975, Chapter 2, pp. 19–22.
5. John Gaschnig, *Performance Measurements and Analysis of Certain Search Algorithms*, Ph.D. Thesis, Department of Computer Science, Carnegie-Mellon University, May 1979.
6. Thomas C. Henderson and Ashok Samal, "Multi-constraint shape analysis", *Image and Vision Computing* **4** (1986) 84–96.
7. Ugo Montanari, "Networks of constraints: fundamental properties and applications to picture processing" *Information Sciences* **7** (1974) 95–132.
8. Donald E. Knuth, "Estimating the efficiency of backtrack programs", *Mathematics of Computation* **29** (1975) 121–136.
9. Robert M. Haralick and Gordon Elliot, "Increasing Tree Search Efficiency for Constraint Satisfaction Problems", Technical Report, Virginia Polytechnic Institute and State University, March 1979.

10. Robert M. Haralick, Larry S. Davis, Azriel Rosenfield and David Milgram, "Reduction operations for constraint satisfaction", *Information Sciences* **14** (1978) 199–219.
11. Robert M. Haralick and Linda G. Shapiro, "The consistent labeling problem: Part II", *IEEE Transactions On Pattern Analysis And Machine Intelligence* **2** (1980) 193–203.
12. Azriel Rosenfield, Robert A. Hummel and Steven W. Zucker, "Scene labeling by relaxation operations", *IEEE Transactions On Systems, Man And Cybernetics* **6** (1976) 420–433.
13. Roger Mohr and Thomas C. Henderson, "Arc and path consistency revisited", *Artificial Intelligence* **28** (1986) 225–233.
14. John Gaschnig, "A constraint satisfaction method for inference making" *Proc. 12th Annual Allerton Conf. Circuit and Sysem Theory*, U. Ill. Urbana-Champaign, 1974, pp. 866–874.
15. Larry S. Davis and Thomas C. Henderson, "Hierarchical constraint processes for shape analysis", *IEEE Transactions On Pattern Analysis And Machine Intelligence* **3** (1981) 265–277.
16 Thomas C. Henderson and Larry Davis, "Hierarchical models and analysis of shape", *Pattern Recognition* **14** (1981) 197–204.
17. D. Nitzan, et al., "Machine Intelligence Research Applied to Industrial Automation", Technical Report, SRI International, Jan. 1982.
18. Tom Henderson and Ashok Samal, "2-D scene analysis using split-level relaxation", *Proc. of the 8th International Conference on Pattern Recognition*, Oct. 1986, pp. 195–197.
19. Charles L. Seitz, "The Cosmic Cube", *Communications of the ACM* **28** (1985) 22–33.
20. Sequent Computer Systems, Inc., *BALANCE 8000 Guide to Parallel Programming* (1985).
21. Peter Tinker, *The Design and Implementation of an OR-Parallel Logic Programming System*, Ph.D. Thesis, Department of Computer Science, University of Utah, 1987.
22. BBN Laboratories Inc., *Butterfly Parallel Processor Overview*, Report No. 6148, March 6, 1986.
23. BBN Laboratories Inc., *The Butterfly RAMFile System*, Report No. 6351, Sept. 10, 1986.
24. BBN Laboratories Inc., *Chrysalis Programmers Manual Version 2.3.1*, BBN Report No. 6191, Aug. 17, 1986.
25. Ashok Samal, "Parallel discrete relaxation algorithms", in *Discrete Relaxation Techniques* by Thomas C. Henderson, Oxford Press, 1988, Chapter 8. To appear.
26. Ashok Samal and Thomas C. Henderson, "Parallel consistent labeling algorithms", to appear in *International Journal of Parallel Programming*.
27. Ashok Samal, *"Parallel Split-Level Relaxation"*, Ph.D. Thesis, Department of Computer Science, University of Utah, Aug. 1988.