

Washington University in St. Louis
Washington University Open Scholarship

All Computer Science and Engineering Research

Computer Science and Engineering

Report Number: WUCS-90-15

1990-04-01

A Parallel Algorithm for Incremental Stereo Matching on SIMD Machines

Authors: Andrew F. Laine and Gruia-Catalin Roman

Previous matching algorithms have achieved high speeds through algorithm simplification and/or relied on custom hardware. The objective of our work has been the development a robust high-speed stereo matcher by exploiting parallel algorithms executing on general purpose SIMD machines. Our approach is based on several existing techniques dealing with the classification and evaluation of matches, the application of ordering constraints, and relaxation-based matching. The techniques have been integrated and reformulated in terms of parallel execution on the theoretical SIMD machine. Feasibility is demonstrated by implementation on a commercially available SIMD machine. Its performance is compared with that of the idealized machine.

Follow this and additional works at: http://openscholarship.wustl.edu/cse_research

Recommended Citation

Laine, Andrew F. and Roman, Gruia-Catalin, "A Parallel Algorithm for Incremental Stereo Matching on SIMD Machines" Report Number: WUCS-90-15 (1990). *All Computer Science and Engineering Research*.
http://openscholarship.wustl.edu/cse_research/690

**A PARALLEL ALGORITHM FOR
INCREMENTAL STEREO MATCHING ON
SIMD MACHINES**

Andrew F. Laine and Gruia-Catalin Roman

WUCS-90-15

April 1990

**Department of Computer Science
Washington University
Campus Box 1045
One Brookings Drive
Saint Louis, MO 63130-4899**

This paper has been published in the *Proceedings of the 10th International Conference on Pattern Recognition: Computer Architectures for Vision and Pattern Recognition*, and is an abridged version of Technical Report WUCS-89-16.

A PARALLEL ALGORITHM FOR INCREMENTAL STEREO MATCHING ON SIMD MACHINES

Andrew F. Laine and Gruia-Catalin Roman

Department of Computer Science
WASHINGTON UNIVERSITY
Saint Louis, Missouri 63130

ABSTRACT

Previous matching algorithms have achieved high speeds through algorithm simplification and/or relied on custom hardware. The objective of our work has been the development a robust high-speed stereo matcher by exploiting parallel algorithms executing on general purpose SIMD machines. Our approach is based on several existing techniques dealing with the classification and evaluation of matches, the application of ordering constraints, and relaxation-based matching. The techniques have been integrated and reformulated in terms of parallel execution on a theoretical SIMD machine. Feasibility is demonstrated by implementation on a commercially available SIMD machine. Its performance is compared with that of the idealized machine.

1. INTRODUCTION

The goal of stereo vision is the recovery of depth information from the relative lateral displacements in the positions of objects within a pair of images taken from slightly differing viewpoints. The recovered depth information may then be used to reconstruct the three-dimensional structure of a scene. The fundamental problem in stereo vision is the matching of corresponding points in the different views. Because of possible occlusions, not all points may have matches and context information is used to infer the depth of the unmatched points [1]. Techniques differ in the strategy they follow with regard to the generation of a unique and consistent set of matches.

Recent techniques have been able to reduce the search space required to maintain global consistency between matches and have yielded speed improvements without compromising reliability [2]. However, these algorithms still fall short of the real-time stereo matching requirements for navigation systems, robot vision, machine inspection and other areas of computer vision where rapid response is critical.

Some matching algorithms have achieved high speeds through algorithm simplification [3] and/or relied on custom hardware [4]. The objective of our work has been the development of a robust real-time stereo matcher by exploiting parallel algorithms executing on general purpose SIMD machines. Our approach is based on several existing techniques dealing with the classification and evaluation of matches [5], the application of ordering constraints [6], and relaxation-based matching [7]. The techniques have been integrated and reformulated in terms of parallel execution on a theoretical SIMD machine. The algorithm was then implemented on a commercially available SIMD machine.

Our parallel algorithm consists of two phases. First, unlikely matches are discarded based on a loose geometric constraint and the ordering of any previous matches. Remaining matches are then evaluated using criteria based on precomputed similarity measures (such as direction and intensity on each side of an edge). Each set of matches, perhaps containing several

candidate matches, is classified based on the aggregate of the previously evaluated candidate matches. Finally, each set of matches is sorted and truncated so that it contains no more than three of the most likely candidates.

Phase two computes initial estimates of the probability of each possible match based on the individual evaluation of the match and the classification of its set. These initial estimates are refined during a relaxation process, by a consistency rule that successively increases (decreases) the probability of matches if nearby points have similar (different) disparity. Afterwards, accepted matches are identified and the entire algorithm may be repeated. At the start of each iteration, previously accepted matches may be used to provide a context for the incremental accumulation of new matches.

The machine model on which the parallel algorithm has been formulated assumes a 2-D array of pipelined processors and a set of memory arrays that may be read and/or updated during each machine cycle. Each processor is capable of performing four kinds of operations: logical, integer arithmetic, max/min, and functions of one variable. Model parameters include the number of stages per pipeline, input and output bandwidth and stage interconnection bandwidth.

The performance of the parallel algorithm is directly related to maximum disparity (δ) and can run in constant time on an idealized machine having 2δ pipelines, four stages per pipeline, 16-bit interconnection links and an output bus no wider than six pipelines, (6×16) -bits. The parallel algorithm has been implemented as part of an interactive environment [3], consisting of a stereo workstation driven by dual Gould/DeAnza IP-8500's, and a MicroVax II host. Performance measurements show that on a typical commercial pipelined SIMD machine the algorithm accomplishes the match of a $512 \times 512 \times 8$ -bit pair of stereo images with a maximum disparity range (δ) of 32 pixels in less than 30 seconds.

The remaining sections of this paper are organized as follows. Section 2 presents an informal overview of the matching strategy. Section 3 provides a formal description of the SIMD machine model. Section 4 contains a detailed description of the parallel algorithm and the resources it requires. Section 5 compares the performance analysis for a theoretically ideal machine and an actual implementation. Finally, a discussion and conclusions are presented in section 6.

2. MATCHING STRATEGY

This section presents an overview of the matching process including all necessary preprocessing steps. We assume an epipolar camera model (i.e., the horizontal scan lines of both cameras are parallel to the baseline so that all disparities are horizontal) and constrain searches for candidate matches to some predetermined disparity range. While the primitive features matched are edgels (edge pixels), the similarity measures used to compare features are based on the properties of edge segments (i.e., chains of edgels). The paragraphs below describe

the preprocessing steps, candidate selection criteria, and relaxation method used to produce a consistent set of matched features.

Preprocessing. Edgels with magnitude and direction information are produced by a Kirsh edge detector [8]. Edge thinning is accomplished using nonmaxima absorption in the gradient direction [9]. The edge magnitudes are then thresholded and the resulting binary edgels thinned, producing 8-connected edge segments one pixel wide. Non-horizontal chains have at most one pixel per scanline when the images are aligned along the epipolar axis; horizontal and nearly horizontal chains have at most one pixel per column.

Next, for each chain in the image we compute its direction and the average intensity along its left and right sides. The precomputed properties are stored as *characteristic images*. A characteristic image associates with each edgel in a chain a value in the range [1-255] corresponding to some property of the whole chain to which the edgel belongs. Image pixels that do not belong to a chain are assigned the value zero.

Candidate selection. Input to the matcher consists of a set of selected edgels, three sets of characteristic images and, (optionally) a set of previously matched edgels. Intervals between any previously matched edgels are labeled sequentially to provide a global context in which to embed new matches. Previously matched edgels are then removed from the set of selected edgels to avoid the possibility of rematching. An initial set of possible matches is constructed by pairing each edgel in the left image with every edgel from the right image which is on the same epipolar line and within some distance δ of the location of the edgel in the left image. Some of these candidates may be eliminated by the ordering constraint imposed by previous matches [3], and the orientation of corresponding chains.

The remaining candidates within the pool are sorted by evaluating the degree of similarity between the precomputed properties of corresponding chains. The evaluation process results in partitioning the pool of candidates into three disjoint sets with respect to similarities in orientation, left-side and right-side intensity.

Next, the pool itself is classified based on the aggregate of its previously evaluated candidates. Pools containing a single (unique) candidate or a candidate that is discernibly better than all other candidates are preferred. This classification information is used later to assign initial estimates of probabilities of candidates for relaxation.

Following the classification process, each pool is truncated so that it contains no more than three of the most promising candidates. Since more than one promising candidates are kept, the opportunity to reevaluate their credibility and to correct some false matches remains. We do this by employing a relaxation method that corrects most local errors [10] by relying on two types of continuity constraints.

Relaxation. Initial probabilities are assigned to each candidate in the pool, using a simple weighting function which takes into account the candidate's previous evaluation, and the classification of its pool. These initial estimates are iteratively refined by applying a consistency rule to all the candidates within the pool, updating the probability of each candidate, and normalizing the new probability estimates. The relaxation follows the procedure of Bernard [7] and has been formulated to apply constraints on *figural continuity* as well as *disparity continuity*. Candidates that are continuously connected along the

same figural contour are allowed to support each other in a cooperative sense. After several iterations, consistent candidates *increase* in probability. Conversely, inconsistent candidates, having no *connected* neighbors satisfying the disparity constraint, receive little support, and *decrease* in probability.

In practice, only small changes in probabilities are observed after a few iterations. For this reason, matches with probabilities greater than an empirically selected threshold (0.7) are accepted. Accepted matches are assigned probability 1.0, and saved. In subsequent iterations, these matches provide global support for new candidates that are locally consistent.

The same matching process is also applied to matching features in the right image with features in the left image and a final set of matches is produced by selecting those pairs of edgels that survived both right-to-left and left-to-right matching processes.

3. MACHINE MODEL

The method above has been cast as a parallel algorithm executing on a theoretical pipelined SIMD machine. As shown in Figure 1, the machine consists of a two-dimensional array of processors and a set of dual ported memories. The machine is organized as a system of n pipelines, each consisting of m stages. Stages are linked to form a rectangular network of horizontal pipelines, with no feedback between stages. Data may flow from processor (i, j) , the j th stage of pipeline i , to processors $(i, j+1)$, $(i+1, j+1)$, or $(i-1, j+1)$. Conversely, processor (i, j) may accept inputs from processors $(i, j-1)$, $(i+1, j-1)$, and $(i-1, j-1)$. Input data, supplied from as many as w memories ($w \leq 2n$) are processed and delivered to as many as p memories ($p \leq n$) within a single machine cycle. Each memory may be connected to several pipeline inputs, and read simultaneously by the first stage of each pipeline. However, each memory may be connected to one pipeline output at most. A pipeline may deliver its data to more than one memory within the same machine cycle.




Figure 1. Machine model of a parallel synchronous pipelined network.

For b -bit processors, all links (input, stage, output) are b -bits wide, the input bus has width less than or equal to $2nb$ -bits and the output bus has width less than or equal to pb -bits. Internal processing within each stage includes the following simple operations: logical operations (\wedge , \vee , \neg , \otimes), integer arithmetic

(+, −, ×, min, max), accumulation, and functions of one variable. Each processor may have no more than two forward input links¹ and one output link, active during any one machine cycle. We adopt the notation $A^{[1:N]}$ to denote a group of *contiguous* memories 1 through N . Each distinct memory A^a may consist of one or more contiguous bit-planes, depending on the data type stored. Thus, A^k identifies the k -th memory of group A . When the name of the memory group A stands alone, it represents the complete set of contiguous memories $[1:N]$.

For purposes of making more obvious the mapping of the algorithmic steps to the architecture, operations along the stages of each pipeline are contained between the delimiters "(" and ")", concurrent processing of more than one pipeline is specified by the symbol \parallel , and memory coordinates ranging over an entire image memory are preceded by the symbol \square . For example, consider the statement below where memories A^k and B contain boolean data, one-bit deep:

$$\parallel k: 1 \leq k \leq N :: \langle \square i, j :: A^k(i, j) := \langle A^k(i, j) \vee B(i+k, j) \rangle - B(i+k, j) \rangle.$$

Each (distinct) A^k is computed over a sub-network of two pipelines, each having two stages, executing in parallel and synchronously. The first stage of one pipeline forms the logical union of memories A^k and B , with memory B shifted k pixels along the horizontal axis. The result is passed to the second stage of an adjacent pipeline where B is subtracted from the union, and its output delivered back to memory A^k .

We need two pipelines because each stage may process no more than two inputs per operation. One way of linking the processors needed to carry out this computation is shown above in Figure 1, for the k -th network. The complete computation requires N identical pipelined sub-networks executing in parallel. For a real machine with n pipelines, where $n < 2N$, the operation requires $\lceil 2N/n \rceil$ passes. Memory $B(i+k)$ is read by the first stages of each distinct sub-network simultaneously. All the pixels within a memory ($\square i, j$) are processed uniformly within each pipeline in a single machine cycle. Thus, all A^k 's are updated in parallel in one machine cycle.

The notation serves only to help the reader in understanding the parallel algorithm presented in this paper, and is not intended to be a complete specification language for parallel synchronous computations in general. Each statement consists of a quantified expression of the form $\langle \text{operator variables} : \text{range} :: \text{expression} \rangle$. Allowed operators are $\parallel, \forall, \square, +, -, \times, \vee, \wedge, \min$ and \max . An expression may be an assignment, a conditional assignment, or another quantified expression. A statement is evaluated by applying an *operator* to the set of *expressions* obtained by instantiation over the *range* of bound *variables*. If the range is obvious it is omitted. For example, $\square i, j \equiv \forall i, j : x_{\min} \leq i \leq x_{\max}, y_{\min} \leq j \leq y_{\max}$. In the next section we use this notation to present a parallel version of the matching algorithm described earlier.

4. ALGORITHM

In describing the algorithm, we will use L and R to refer to the sets of edgels belonging to chains selected from the left and right images. M_L and M_R will refer to any edgels previously matched. These sets are encoded as binary two-dimensional arrays with the index ranging over the image space. $L(i, j)$, for instance, assumes the value "1" if the position (i, j) contains an edgel belonging to a chain selected from the left image and the

value "0" otherwise. As described earlier, the first phase of the matching algorithm identifies the three most promising candidates within each match pool. The second phase applies a relaxation method to identify the most consistent of the three candidates.

Phase 1. For sake of clarity, we restrict our discussion to matching features from the left image to features in the right image. Corresponding right-based matching may be performed in parallel, and is symmetric with respect to the equations and steps below. The first phase is described in nine steps.

(Step 0) Label intervals between previous matches. The set of previously matched edgel-pairs is a bijection. Therefore on each epipolar line, the intervals between matched edgels match pairwise. In this step, we label the intervals between already matched edgels; we assign identical labels to all pixels within a given interval. The labels are used later to enforce an ordering constraint.

$$\square i, j :: B_L(i, j) := (+\text{mod}(2^b - 1))y, x : (0 < y \leq j) \wedge (0 < x \leq i) :: M_L(x, y)$$

$$\square i, j :: B_R(i, j) := (+\text{mod}(2^b - 1))y, x : (0 < y \leq j) \wedge (0 < x \leq i) :: M_R(x, y)$$

$B_L(i, j)$ is computed by counting the number of previously matched edgels from $M_L(0, 0)$ to $M_L(i, j)$. The result is that all pixels within matching intervals have the same label. All the labels stored in B_L and B_R are computed in parallel in one machine cycle. Memories B_L and B_R are at least b -bits wide, such that $(2^b - 1)$ is at least the length of a scanline. Thus label values remain unique within each scanline.

(Step 1) Eliminate previous matches. Any previous matches (M_L and M_R) are eliminated from the set of selected edgels (L and R). Memories M_L and M_R supply input data for both pipelines simultaneously, allowing L and R to be modified in parallel.

$$\square i, j :: \langle L(i, j) := \langle L(i, j) \vee M_L(i, j) \rangle - M_L(i, j) \rangle$$

$$\square i, j :: \langle R(i, j) := \langle R(i, j) \vee M_R(i, j) \rangle - M_R(i, j) \rangle$$

(Step 2) Build disparity map. For each edgel (i, j) in the left image and disparity k in the range 0 to δ , if an edgel exists in the right image a bit is set in an image memory CP_L^k . For each match pool associated with an edgel in L , candidates from R having disparity k are identified by computing the conjunction of input edgel sets L and R , where R has been shifted k pixels along the epipolar axis.

$$\parallel k : 0 < k < \delta :: \langle \square i, j :: CP_L^k(i, j) := L(i, j) \wedge R((i+k), j) \rangle$$

All candidates are identified in a single machine cycle by a stack of δ independent pipelines executing in parallel.

(Step 3) Apply ordering constraint. In this step, candidates are eliminated from further consideration if they do not fall within matching intervals, i.e., if they do not have the same label number. Candidates not violating the ordering constraint simply pass through and are returned to memory CP_L for further processing.

$$\parallel k : 0 < k < \delta :: \langle \square i, j :: CP_L^k(i, j) := \text{if } \langle B_L(i, j) - B_R((i+k), j) = 0 \rangle \text{ then } CP_L^k(i, j) \text{ else } 0 \rangle$$

Candidates for all δ disparities, each stored in a distinct memory CP_L^k , are processed in parallel by δ pipelined networks. Thus, candidates of all disparities are tested in one machine cycle.

¹However, in practice, lateral inputs for bit-sliced arithmetic and carry operations are allowed when necessary.

(Step 4) *Apply orientation constraint.* L_{DIR} and R_{DIR} are characteristic images containing the precomputed values for orientation of chains belonging to the left and right edge sets respectively. The absolute difference between the orientations of corresponding chains belonging to a candidate is compared to a preselected threshold α . For candidates of disparity k , property values in the right characteristic image are found by shifting memory R_{DIR} k pixels along the epipolar axis. Candidates of all δ disparities are compared in parallel. If for some candidate $CP_L^k(i,j)$, the difference is less the threshold (e.g., $\alpha=30$), the candidate remains viable, and is stored in $CP_L^k(i,j)$. Otherwise it is discarded.

$$\|k: 0 < k < \delta :: \langle \square i,j :: CP_L^k(i,j) \rangle :=$$

$$\text{if } \langle |L_{DIR}(i,j) - R_{DIR}((i+k),j)| < \alpha \rangle \text{ then } CP_L^k(i,j) \text{ else } 0 \rangle$$

The computational structure is similar to the ordering constraint described above, and requires only a single machine cycle to process candidates at all δ disparities.

(Step 5) *Count candidates within each pool.* For each candidate pool $CP_L^{1:\delta}(i,j)$, the surviving candidates are counted and the total sum is stored in memory $COUNT_L$. Candidates at all δ possible disparities are tallied in parallel by one or more function tables (f_{count}) that count the number of bits set within each $CP_L(i,j)$ pool. Since it is only possible to count δ/b -bits² at a time, if $\lceil \frac{\delta}{b} \rceil \geq 1$, memories $CP^{1:\delta}$ are partitioned into $\lceil \frac{\delta}{b} \rceil$ disjoint groups of contiguous bitplanes. The k -th distinct group is addressed below as memory $CP^{[k]}$. The sums for all candidate pools are computed in a single machine cycle.

$$\langle \square i,j :: COUNT_L(i,j) := (+k : 1 \leq k \leq \lceil \frac{\delta}{b} \rceil :: f_{count}(CP_L^{[k]}(i,j))) \rangle^2$$

(Step 6) *Identify candidates satisfying at least one similarity constraint.* The three parallel computations below inject the candidates of CP_L into three *similarity groups*, S_{DIR} , S_{LS} and S_{RS} . (L_{LS} , R_{LS}) and (L_{RS} , R_{RS}) are pairs of characteristic images for the average intensity along the left side and right side of chains within edge images L and R , respectively. After the completion of one machine cycle, memory S_{DIR} contains all the candidates whose corresponding chains are similar in orientation. For example, consider a candidate identified by the bit set in memory $CP_L^k(i,j)$; the candidate (bit) will be copied into memory $S_{DIR}^k(i,j)$ if its corresponding property values $L_{DIR}(i,j)$ and $R_{DIR}(i+k,j)$ differ (absolutely) by less than $max.dir$.

$$\|k: 0 < k < \delta :: \langle \square i,j :: S_{DIR}^k(i,j) \rangle :=$$

$$\text{if } \langle |L_{DIR}(i,j) - R_{DIR}((i+k),j)| < max.dir \rangle \text{ then } CP_L^k(i,j) \text{ else } 0 \rangle$$

Likewise, memories S_{LS} and S_{RS} contain candidates whose corresponding chains are similar in left-sided and right-sided intensity, respectively. Candidates may be included in more than one S -group.

$$\|k: 0 < k < \delta :: \langle \square i,j :: S_{LS}^k(i,j) \rangle :=$$

$$\text{if } \langle |L_{LS}(i,j) - R_{LS}((i+k),j)| < max.ls \rangle \text{ then } CP_L^k(i,j) \text{ else } 0 \rangle$$

$$\|k: 0 < k < \delta :: \langle \square i,j :: S_{RS}^k(i,j) \rangle :=$$

$$\text{if } \langle |L_{RS}(i,j) - R_{RS}((i+k),j)| < max.rs \rangle \text{ then } CP_L^k(i,j) \text{ else } 0 \rangle$$

All three parallel computations are similar to Step 4 in structure. Candidates at all δ disparities are identified in parallel in one machine cycle. The three similarity groups S_{DIR} , S_{LS} and S_{RS} are used in the next step to evaluate all candidates in parallel.

(Step 7) *Evaluate each candidate within a pool.* Each candidate within a pool (i,j) is assigned to one of three disjoint partitions based on "goodness" of match. The highest quality candidates are assigned to partition T_1 and are identified by the conjunction of candidates belonging to similarity groups S_{DIR} , S_{LS} and S_{RS} . The conjunction of candidates of all δ disparities is computed in parallel.

$$\langle \square i,j :: T_1^{1:\delta}(i,j) := \langle S_{DIR}^{1:\delta}(i,j) \wedge S_{LS}^{1:\delta}(i,j) \wedge S_{RS}^{1:\delta}(i,j) \rangle$$

Candidates are assigned to partition T_2 by first computing the disjunction of candidates in similarity groups S_{RS} and S_{LS} . Candidates of the union are then intersected with candidates of similarity group S_{DIR} . To keep partitions T_2 and T_1 disjoint, candidates must not have already been assigned to partition T_1 .

$$\langle \square i,j :: T_2^{1:\delta}(i,j) :=$$

$$\langle S_{DIR}^{1:\delta}(i,j) \wedge \langle S_{LS}^{1:\delta}(i,j) \vee S_{RS}^{1:\delta}(i,j) \rangle \rangle \wedge \neg T_1^{1:\delta}(i,j)$$

Candidates assigned to the third partition, T_3 , are simply the remaining candidates that have not been assigned to either partition T_1 or T_2 .

$$\langle \square i,j :: T_3^{1:\delta}(i,j) := CP_L^{1:\delta}(i,j) \otimes \langle S_{LS}^{1:\delta}(i,j) \vee S_{RS}^{1:\delta}(i,j) \rangle \rangle$$

$COUNT_{T_1}$ and $COUNT_{T_2}$ are computed similarly to Step 5 above, and are used in the next step to classify each match pool in parallel.

$$\langle \square i,j :: COUNT_{T_1}(i,j) := (+k : 1 \leq k \leq \lceil \frac{\delta}{b} \rceil :: f_{count}(T_1^{[k]}(i,j))) \rangle$$

$$\langle \square i,j :: COUNT_{T_2}(i,j) := (+k : 1 \leq k \leq \lceil \frac{\delta}{b} \rceil :: f_{count}(T_2^{[k]}(i,j))) \rangle$$

(Step 8) *Classify each pool into one of four disjoint sets.* In each classification operation below, all pools are processed in parallel. A pool at position (i,j) having exactly one candidate at some disparity k belonging to partition T_1 is identified in memory $CLASS_{U1}$ (Unique of type I) at the same position. Similarly, pools in image memory $CLASS_{U2}$ have a unique candidate belonging to partition T_2 . The pools identified in image memory $CLASS_{B1}$ have more than one candidate, but exactly one of them belongs to partition T_1 , making it discernibly *better* than the rest. Similarly, the pools of $CLASS_{B2}$ have more than one candidate, but have exactly one candidate belonging to T_2 and exactly zero candidates belonging to partition T_1 .

$$\langle \square i,j :: CLASS_{U1}(i,j) := \langle COUNT_{T_1}(i,j) = 1 \wedge COUNT_L(i,j) = 1 \rangle$$

$$\langle \square i,j :: CLASS_{U2}(i,j) := \langle COUNT_{T_2}(i,j) = 1 \wedge COUNT_L(i,j) = 1 \rangle$$

² In this case, our notation fails to explicitly specify the linkage between the stages of the pipes used to accomplish the (binary) addition operation. However, for N inputs, the operation requires at most $\lceil \frac{(N/2)+1}{2} \rceil$ stages per pipe.

$$\langle \square i,j :: CLASS_{B1}(i,j) := (COUNT_{T_1}(i,j)=1 \wedge COUNT_L(i,j) > 1) \rangle$$

$$\langle \square i,j :: CLASS_{B2}(i,j) := (COUNT_{T_1}(i,j)=0) \wedge (COUNT_{T_2}(i,j)=1) \wedge (COUNT_L(i,j) > 1) \rangle$$

The four classes above may be computed in a single machine cycle by four independent pipelines executing in parallel. The four classes are consolidated into a single memory C by the disjunction of all classes. Each distinct function f_c outputs a unique value labeling $CLASS_c$ pools.

$$\langle \square i,j :: C(i,j) := (\vee c : c \in \{U1,U2,B1,B2\} :: f_c(CLASS_c)) \rangle$$

Phase 2. The next phase of the computation consists of a relaxation method that employs constraints on figural continuity [11,12] and disparity continuity to correct most local errors. Below, phase two is described in six steps.

(Step 9) *Identify the three most likely candidates from each match pool.* As described earlier in step seven, we partitioned all candidates into one of three disjoint sets, T_1, T_2 and T_3 . In the equation below, the variable t ranges over the number Typ of candidate types, so in our method, $Typ=3$. In order to identify the most promising candidates first, binding of the variable t proceeds from 1 (the most promising candidates are of type T_1) to 3 (the least promising candidates are of type T_3). The variable n ranges over the number N of (best) candidates saved. In our method, we save the three most likely candidates within each pool, so $N=3$. For each t , as many as N distinct candidates are identified in parallel. At the completion of this step, memories $D^n(i,j)$ and $W^n(i,j)$ will respectively contain disparity and weight values of the n -th candidate of pool (i,j) . Initially all N disparity memories (D^n) and weight memories (W^n) are cleared. If for some type t , there exists a candidate within pool (i,j) of T_t (and $D^n(i,j)$ is empty) then memories $D^n(i,j)$ and $W^n(i,j)$ are assigned disparity and weight values of the candidate, respectively.

$$t = 1..3 :: |n : 1 \leq n \leq N ::$$

$$\langle \square i,j :: D^n(i,j), W^n(i,j) := \\ \text{if } \langle (D^n(i,j)=0) \wedge (f_{type}^n(T_t^{1:\delta}(i,j)) \neq 0) \rangle \\ \text{then } f_{disparity}^n(T_t^{1:\delta}(i,j)), f_{weight}^n(C(i,j)) \\ \text{else } D^n(i,j), W^n(i,j) \rangle$$

Let us examine how all N pairs of disparity and weight values are computed in parallel, for each candidate type t . Consider a pool (i,j) . The function f_{type}^n takes as input all the candidates of pool (i,j) stored in memory T_t and returns the value "1" if there exists at least n candidates in the pool. Similarly, the function $f_{disparity}^n$ takes as input all the candidates of pool (i,j) and assigns a disparity value to (exactly) the n -th candidate of each pool, relative to the (bit-plane) index $[1:\delta]$ of memory T_t . For example, if the n -th candidate of a pool is stored in memory T_t^7 then it (the n -th candidate) is assigned a relative disparity value "7". The function f_{weight}^n assigns a weight to each admitted candidate by a table lookup of precomputed values³. If during a subsequent iteration t , there exists

³ Each value is a measured sum which takes into account the type of candidate and the classification of its pool. Weights range from 0.0 (the least promising candidates) to 1.0 (the most promising candidates). There are Typ distinct f_{weight}^n functions precomputed as follows:

another candidate within the same (i,j) pool, and disparity memory D^n contains a previous entry for pool (i,j) , the disparity entry and weight value are not effected by the pass. Similarly, if there is no candidate within pool $T_t(i,j)$ and disparity $D^n(i,j)$ is empty, memories D^n and W^n are left unchanged.

The best candidates of type t (at any of δ possible disparities) are identified in parallel. After Typ machine cycles, memory groups $D^{1:N}$ and $W^{1:N}$ contain the disparities and associated weights for the N most likely candidates, respectively. The weight of each candidate is used in the next step to compute probability values for all N candidates of each pool in parallel.

(Step 10) *Calculate the initial probability of each candidate in every pool.* This computation is accomplished in three steps. First, the maximum weight of the N candidates within each pool is identified in parallel. Image memory I^* contains an initial estimate of the probability that every pool (i,j) is matchable.

$$\langle \square i,j :: I^*(i,j) := \langle 1.0 - (\max n:1 \leq n \leq N :: W^n(i,j)) \rangle \rangle$$

Next, the sum of all N candidates within each pool (i,j) is computed and stored in $S(i,j)$.

$$\langle \square i,j :: S(i,j) := (+ n:1 \leq n \leq N :: W^n(i,j)) \rangle$$

Each candidate weight of pool (i,j) is converted into an initial probability estimate through normalization. The function $f_{reciprocal}$ is a precomputed table that outputs the inverse of it's input value.

We use fixed point arithmetic for all analytic computations.⁴ Fixed point division is accomplished by computing the reciprocal (at one stage) and then multiplying (at a subsequent stage). We can avoid costly real division without the loss of precision because we can determine the domain of input values for any pool size N . For example, since each candidate within a given pool is assigned a measured weight less than or equal to 1.0, the total sum of all candidates is less than or equal to N , and the range of $S(i,j)$ is constrained to the interval $[0:N]$.

$$| | n : 1 \leq n \leq N :: \langle \square i,j :: P^n(i,j) :=$$

$$\langle W^n(i,j) \times f_{reciprocal}(S(i,j)) \rangle \times \langle 1 - I^*(i,j) \rangle \rangle$$

All N candidates of every pool are normalized in parallel, at a cost of one machine cycle. At the end of the cycle, the memories $P^{1:N}$ contain the normalized probability values for all N candidates of every pool. In the next step, these normalized probabilities are used to determine how consistent each candidate is with it's neighbors.

(Step 11) *Compute local support for each candidate.* The local support of each candidate is computed by summing the probability values of connected neighbors. Only connected neighbors (above or below) with disparity differences less than one pixel are allowed to contribute support. The total support for the n -th candidate of pool (i,j) is stored in memory $Q^n(i,j)$. The variables l and k scroll memories D^m and P^m in

$$t, c : t \in \{1,2,4\} \wedge c \in \{1,2,8\} :: f_{weight}(c) = 0.6 \left[\frac{1}{1/c} \right] + 0.4 \left[\frac{1}{1/l} \right]$$

Values of the variable t are associated with candidate types 1, 2 and 3, respectively. Values of the variable c are associated with *unique* classes (U1,U2), classes containing a discernibly *better* candidate (B1,B2), and candidates belonging *other* pools (pools with a single type 3 candidate or multiple candidates of the same type), respectively.

⁴In practice, we found 16-bits of precision yielded sufficient accuracy for all numerical computations of this algorithm.

phase with respect to pool (i, j) so that the six connected neighbors (three above and three below) of $D^n(i, j)$ are allowed to contribute support for candidate $D^n(i, j)$, provided they are consistent with candidate $D^n(i, j)$.

$$n: 1 \leq n \leq N :: \langle \square i, j :: Q^n(i, j) :=$$

$$(+k, l, m : k \in \{-1, 0, 1\} \wedge l \in \{-1, 1\} \wedge 1 \leq m \leq N ::$$

$$(f_{\text{consistent}}((D^n(i, j) - D^m(i+k, j+l))) \times P^m(i+k, j+l))) \rangle$$

The function $f_{\text{consistent}}$ returns the value "1" if the (absolute) difference between disparity $D^n(i, j)$ and some connected neighbor disparity $D^m(i+k, j+l)$ is less than or equal to one. Otherwise the function returns the value "0". The output of the function is multiplied by the probability value of the same connected neighbor $P^m(i+k, j+l)$. There are two levels of parallelism taking place. The probabilities of all six connected neighbors are allowed to contribute to the sum of $Q^n(i, j)$ in parallel. The support available from all the neighbors of neighborhood m is gathered simultaneously for all N candidates in each pool. Any previous matches (assigned probability 1.0) provide a network of global support for any new locally consistent matches. The operation is repeated N times (over the variable n) to allow each candidate to obtain support from N distinct neighborhoods. In the next step, the quantity of support stored in memory $Q^n(i, j)$ for each candidate, is used to modify (increase or decrease) in parallel the probability values stored in memory $P^n(i, j)$.

(Step 12) *Update the probability of each candidate.* In the computation below, α and β are parameters that influence the convergence characteristics of the updating rule. Briefly, α and β can be interpreted as damping and gain parameters. The value α delays the suppression of unlikely candidates, and β determines the rate of convergence [13]. In our experiments, $\alpha = 0.3$ and $\beta = 3$.

$$|| n : 1 \leq n \leq N :: \langle \square i, j :: \hat{P}^n(i, j) :=$$

$$\langle \langle P^n(i, j) \times \alpha \rangle + \langle \langle P^n(i, j) \times \beta \rangle \times Q^n(i, j) \rangle \rangle$$

For all N candidates, the probabilities stored in each pool of P^n are multiplied by constants α and β in parallel. Image memories $Q^{1:N}$ and $P^{1:N}$, containing the quantity of local support, and the previous probability of each candidate, are read simultaneously. The resulting modified probability estimates for all N candidates of each pool are summed in parallel and stored in memories $\hat{P}^{1:N}(i, j)$. The next two steps normalize these updated probability estimates in parallel.

(Step 13) *Compute the sum of probability estimates for candidates within each pool.* The sum of the updated probabilities for all N candidates within each pool is computed in parallel and stored in $S(i, j)$.

$$\langle \square i, j :: S(i, j) := \langle (+ n : 1 \leq n \leq N :: \hat{P}^n(i, j)) \rangle$$

(Step 14) *Normalize probability estimates \hat{P} and l^* .* Finally, the probability estimates are normalized in parallel for all N candidates within each pool (i, j) . This step is similar to Step 1 except that memory l^* is updated in parallel through an independent pipeline. The final normalized probabilities are stored in memories $P^{1:N}$.

$$|| n : 1 \leq n \leq N :: \langle \square i, j :: P^n(i, j), l^*(i, j) :=$$

$$\hat{P}^n(i, j) \times f_{\text{reciprocal}}(S(i, j)), l^*(i, j) \times f_{\text{reciprocal}}(S(i, j)) \rangle$$

Steps 11 through 14 are repeated until the network reaches steady state or until only small changes in probabilities are observed. In our experiments, the network required at most four iterations to stabilize. In the next section, we compare the optimal performance of the parallel algorithm with that of a real machine.

5. PERFORMANCE ANALYSIS

In this section we compare the performance of the algorithm for ideal, typical and implementation machines. A formal comparison and complexity analysis may be found in [14]. All of the steps of phase one can be executed in constant time for any disparity range δ , provided the network has at least 2δ pipelines. Similarly, most of the steps of phase two can be executed in constant time with the exception of steps eight and ten. The run time of step eight is directly related to the number of distinct candidate types, Typ , while the run time of step ten is exactly N , the number of *best* candidates saved from each pool. Both parameters are typically small numbers. Assuming $Typ=3$ and $N=3$, minimum computation time is achievable by an idealized machine having the following configuration: 2δ pipelines (n), 21 stages (m), 16 bits/link (b), 6 outputs (p). The cost of the parallel algorithm is a function of the disparity parameter δ alone. Therefore, parallel architectures that can accommodate large input bandwidth (n) are most suited for high-speed stereo matching applications. Only steps eight and eleven required more than four stages to achieve optimal time. This suggests that the topology of such a pipelined network should consist of a large number of pipelines with few stages.

An ideal machine, operating at 60 Hz, can accomplish stereo matching in 1.5 seconds, using 88 machine cycles. However, typical commercial image processing machines (e.g., Pixar, Vicon, Comptel, I²S, and Sun/TACC) often have less resources and run at slower clock speeds. For a machine with four pipelines, four stages per pipeline, and 8-bit interconnection links we could expect to accomplish stereo matching in about 11.9 seconds, using 357 cycles, at 30 Hz. There exists a four fold difference between the number of cycles needed on a theoretical machine and the number of cycles needed for most commercially available machines. Depending on the requirements of the application, this difference may be tolerable for high-speed stereo matching.

We implemented each step of the algorithm on a Gould/DeAnza IP-8500 image processor, equipped with a Digital Video Processor (DVP). The architecture of the DVP consists of a network of four pipelines, with limited interconnection links. Input bandwidth of the network is 8×10 -bits ($b \times 2n$). Output bandwidth is 4×8 -bits ($p \times b$). There are seven distinct stages within each pipeline. The stages of the DVP are not uniform and could not perform the complete set of operations described for the stages of our idealized machine. On account of this mismatch, it was often necessary to adjust the algorithm to accommodate differences in functionality and other considerations (idiosyncrasies) of the DVP architecture. Stereo matching was achieved in 13.5 seconds using 404 DVP cycles. Actual execution time (wall clock) is about 23 seconds due to the hidden cost of down loading at each step pipeline configuration instructions and stage opcodes from a MicroVax II host. In general, the disparity in performance between the ideal machine and the implementation machine can principally be attributed to the limited (specialized) functionality of the stages, lack of flexible (dynamic) interconnection links, and limited (insufficient) input bandwidth.

6. CONCLUSIONS

We have presented a parallel algorithm for stereo matching that achieves high-speed by exploiting the parallel architectures of typical SIMD processors. The cost of the parallel algorithm is a function of maximum disparity (δ) alone, when executing on an idealized machine having a small number of stages, reasonable interconnection bandwidth, and modest output bandwidth. Parallel architectures that can accommodate large input bandwidth are best suited for high-speed stereo matching applications. Only two steps of the algorithm required more than four stages per pipeline to achieve optimal time. This suggests that the topology of a pipelined network tuned for stereo matching should consist of a *large number of pipelines with few stages*. The parallel formulations presented in this paper exhibit a methodology that may be useful to others in the computer vision community who are interested in reformulating existing low-level vision algorithms for high-speed parallel execution.

The feasibility of our parallel algorithm was shown by implementation on a typical commercial SIMD pipelined processor. As advances in hardware technology continue to reduce machine cycle times, and highly interconnected multicomputers with flexible topologies are produced at practical costs, real-time performance of this parallel algorithm will be achievable and appealing for applications where rapid response is critical.

Acknowledgement: This work was supported by the Center for Intelligent Computer Systems and by the Department of Computer Science at Washington University.

7. REFERENCES

- [1] D. Arnold and T. O. Binford, "Geometric constraints in stereo vision: Image processing for missile guidance," in *Proc. Soc. Photo-Opt. Instrum. Engr.*, vol. 238, pp. 281-292, 1980.
- [2] Y. Ohta, K. Takano and K. Ikeda, "A highspeed matching system based on dynamic programming," in *Proc. First International Conference on Computer Vision*, London, England, June 8, 1987, pp. 335-342.
- [3] G.-C. Roman, A. F. Laine and K. C. Cox, "Interactive complexity control and high-speed stereo matching," in *Proc. Computer Society Conference on Computer Vision and Pattern Recognition*, Ann Arbor, Michigan, June 5, 1988, pp. 171-176.
- [4] H. K. Nishihara, "PRISM: A practical realtime imaging stereo matcher," in *Intelligent Robots: Proc. Third International Conference on Robot Vision and Sensory Controls RoViSeC3*, Cambridge, Massachusetts, November 7, 1983, pp. 134-142.
- [5] M. Pietikainen and D. Harwood, "Depth from three camera stereo," in *Proc. Computer Vision and Pattern Recognition Conference*, Miami Beach, Florida, June 22, 1986, pp. 2-8.
- [6] A. L. Yullie and T. Poggio, "A generalized ordering constraint for stereo correspondence," MIT AI Lab Technical Report AIM-777, Cambridge, MA, May 1984.
- [7] S. T. Barnard and W. B. Thompson, "Disparity analysis of images," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-2, no. 4, pp. 333-340, July 1980.
- [8] R. A. Kirsch, "Computer determination of the constituent structure of biological images," *Computers and Biomedical Research* 4, vol. 3, pp. 315-328, June 1971.
- [9] R. B. Eberline, "An iterative gradient edge detection algorithm," *Computer Graphics and Image Processing*, vol. 5, pp. 245-253, 1976.
- [10] R. Mohan, G. Medioni and R. Nevatia, "Stereo Error Detection, Correction, and Evaluation," *IEEE Transactions of Pattern Analysis and Machine Intelligence*, vol. 11, no. 2, February 1989.
- [11] J. E. W. Mayhew and J. P. Frisby, "The computation of binocular edges," *Perception*, vol. 9, pp. 69-86, 1980.
- [12] ---, "Psychophysical and computational studies towards a theory of human stereopsis," *Artificial Intelligence*, vol. 17, 1981.
- [13] A. Rosenfeld, R. A. Hummel and S. W. Zucker, "Scene labeling by relaxation operations," *IEEE Transactions on System, Man and Cyber.*, vol. SMC-6, June 1976.
- [14] A. F. Laine and G.-C. Roman, "A Parallel Algorithm for High-Speed Stereo Matching", Technical Report WUCS-89-16, Department of Computer Science, Washington University, St. Louis, MO, 1989.