PROCEEDINGS

IEEE
International
Conference on
Robotics and
Automation

Sponsored by
IEEE Robotics and Automation Society

May 2 – 6, 1993
Atlanta, Georgia

Volume 1: May 3, 1993

IEEE Computer Society Press
Los Alamitos, California

Washington • Brussels • Tokyo

Parallel VLSI Architectures for Real-Time Kinematics of Redundant Robots

Ian D. Walker and Joseph R. Cavallaro
Dept. of Electrical and Computer Engineering
Rice University, Houston, TX 77251

Abstract

We describe new architectures for the efficient computation of redundant manipulator kinematics (direct and inverse). By calculating the core of the problem in hardware, we can make full use of the redundancy by implementing more complex self-motion algorithms. A key component of our architecture is the calculation in VLSI hardware of the Singular Value Decomposition of the manipulator Jacobian. Recent advances in VLSI have allowed the mapping of complex algorithms to hardware using systolic arrays with advanced computer arithmetic algorithms, such as the coordinate rotation (CORDIC) algorithms. We use CORDIC arithmetic in the novel design of our special-purpose VLSI array, which is used in computation of the Direct Kinematics Solution (DKS), the manipulator Jacobian, as well as the Jacobian Pseudoinverse. Application-specific (subtask-dependent) portions of the inverse kinematics are handled in parallel by a DSP processor which interfaces with the custom hardware and the host machine. The architecture and algorithm development is valid for general redundant manipulators and a wide range of processors currently available and under development commercially.

1 Introduction

The Robot Forward and Inverse Kinematics problems are critical computationally intensive tasks in real time robot control which can benefit from the judicious application of parallel algorithms and architectures [7, 8]. Parallel architectures including systolic arrays are being successfully applied to more basic algorithms in real-time signal processing and image processing, as well as robotics. Previous work on architecture design has centered on robot kinematics, [12, 13, 20, 22] dynamics, and control [10, 21]. Algorithms and architectures for kinematically redundant manipulators have been considered in [3, 15, 19].

The CORDIC, (Coordinate Rotation Digital Computer), algorithms have been used in robotics for inverse kinematics [9, 12] for nonredundant robots and for control [21]. However, it appears that CORDIC has not been utilized in the Direct Kinematics Solution (DKS) and Jacobian computations, where the

many rotations present fertile ground for its ability to handle rotations efficiently.

The overall objectives of our research are to design a custom VLSI CORDIC architecture suitable for redundant manipulator kinematics, to fabricate chips for this architecture, and implement the resulting architecture in our Robotics Laboratory [19]. In this paper, we describe progress in the first stages of this project, the design and implementation of a high performance VLSI array specifically tailored for the robotic application and the combination of the array with a high performance general purpose processor. Digital Signal Processor (DSP) chips present a number of architectural advantages, including high speed arithmetic and communication ports for parallel array construction.

The self-motion capability of redundant manipulator designs [17] allow them to significantly outperform conventional current robots, particularly in unstructured environments (for obstacle avoidance, fault tolerance, etc.). A number of suggested joint space techniques for self motion control are based on a well-known algorithm using the Jacobian pseudoinverse [17], which must be computed iteratively in real time in practise.

Here, we restructure and recast the velocity level kinematics algorithm to exploit the parallelism in the computation. We then introduce efficient special purpose VLSI arrays to compute the DKS, the Jacobian, and its pseudoinverse. The Singular Value Decomposition (SVD) represents an appealing method of computing the pseudoinverse. Our previous work has shown the applicability of high speed algorithms such as the CORDIC algorithms to the SVD of real matrices [2]. Our architecture features a special purpose array of CORDIC processors to compute the SVD of a matrix. The resulting processor design offers the possibility of implementing in real time some of the more complex schemes for redundant robot kinematic control suggested in the literature.

Our architecture is valid for any type of high performance processor, and is thus not restricted to our particular hardware. We are currently using the Texas Instruments TMS320C30 and TMS320C40 floating-point DSP chips. Based upon the performance of the custom chips designed for the CORDIC SVD array, and utilizing maximum parallelism wherever pos-

sible, we estimate that this architecture will compute each iteration for the joint velocities in approximately 400 micro-seconds. We are currently building custom CMOS chips for this architecture, which is described in more detail in the following.

2 Kinematic Algorithms

The problem of inverse kinematics solves the direct (forward) kinematics (DKS) equations

$$\underline{x} = f(\underline{\theta}) \quad (1)$$

for the values of the n joint variables $\underline{\theta} \in \mathbb{R}^n$ given the m end effector position/orientation variables $\underline{x} \in \mathbb{R}^m$.

Two approaches to inverse kinematics are common for nonredundant robots. One method involves solving (1) directly for the joint angles. Conventionally this involves forming the direct kinematics as a product of 4×4 homogeneous transformation matrices $[H_{i-1}^i(\theta_i)]$ where θ_i is the i th component of $\underline{\theta}$,

$$[H_o^n] = [H_o^1(\theta_1)][H_1^2(\theta_2)] \cdots [H_{n-2}^{n-1}(\theta_{n-1})][H_{n-1}^n(\theta_n)] \quad (2)$$

and equating with a (desired) end effector transformation $[T_o^n]$. The inverse kinematics is then computed by manipulation of (2) to produce n independent equations for the θ_i 's. For many conventional nonredundant robots, (2) yields closed form solutions, and a number of elegant approaches to efficient inverse kinematics computation have been proposed (for example [9, 12]).

In our architecture, we adopt the "resolved-rate" approach, solving for the joint velocities required for end effector motion, using the manipulator Jacobian. By differentiation of (1), we obtain the relationship between the joint velocities and end effector velocities

$$\dot{\underline{x}} = [J]\dot{\underline{\theta}} \quad (3)$$

where the $(m \times n)$ matrix $[J]$ is the Jacobian.

Redundant arm inverse kinematics is typically based at the velocity level on the pseudoinverse of $[J]$:

$$\dot{\underline{\theta}} = [J^+(\underline{\theta})]\dot{\underline{x}} + [I - J^+(\underline{\theta})J(\underline{\theta})]\underline{\epsilon} \quad (4)$$

where $[J^+]$ is the pseudoinverse of the Jacobian (for more details, see [17]). In (4), I is the $(n \times n)$ identity matrix, and $\underline{\epsilon}$ is an $(n \times 1)$ column vector whose values may be arbitrarily selected [17]. Redundant arms have n larger than m (extra joints), and different choices of $\underline{\epsilon}$ determine different possible arm postures, each of which will allow the specified end effector motion $\dot{\underline{x}}$. This capability is however gained at the expense of significantly increased complexity via the need to compute the pseudoinverse and the second term in (4). Earlier we proposed an architecture which computes both the direct kinematics solution (2) and the manipulator Jacobian using a parallel array. The ar-

Figure 1: DSP - CORDIC array configuration.

chitecture then computes the joint velocities using (4) to complete the kinematics [19].

3 VLSI CORDIC / DSP Architecture

We propose an Inverse Kinematics Engine (IKE) architecture, which incorporates a special-purpose VLSI array and a programmable DSP processor. The IKE reduces the computational load on the Real Time Controller. Figure 1 shows the enhanced VLSI CORDIC/DSP architecture demonstrating the Inverse Kinematic Engine (IKE) configuration. The interconnection of the hardware components in the proposed IKE is shown. The CORDIC SVD array acts as a co-processor for the general-purpose DSP processor as shown in Figure 1. In a particular implementation, the Processor Element could be a DSP chip, such as the Texas Instruments TMS320C30 or TMS320C40 [4], or one of the INMOS Transputer chips.

In our architecture, both the DKS and the Jacobian are calculated using the CORDIC array and the DSP chip. To perform the inverse kinematics, the pseudoinverse of the Jacobian is performed in hardware using the CORDIC SVD array. The pseudoinverse is used to calculate the next iteration of joint velocities, with the necessary matrix-matrix and matrix-vector computations carried out by the DSP chip using the information from the CORDIC SVD array. As part of our architecture, we can calculate the DKS in parallel with part of the Inverse Kinematics.

Our architecture has similarities with that of [21] in which a general purpose robotics processor with two CORDIC subunits was devised. However, our special-purpose CORDIC array with multiple CORDIC units

allows us more flexibility in partitioning the various portions of the algorithm in parallel. Our architecture is also more general than that proposed in [5, 10], which is specific to nonredundant robots, and [15], which does not have the CORDIC SVD array or efficient parallel Jacobian calculation. Our approach is valid for more general robots than that in [12], which used a pipelined CORDIC architecture for closed form nonredundant robot inverse kinematics.

4 Parallel DKS and Jacobian Calculations using CORDIC

We use our CORDIC architecture to take advantage of both the parallelism and the large number of rotations inherent in the DKS and Jacobian calculations. We adopt the conventional Jacobian formulation expressed in a base coordinate frame. Other work [16] has considered Jacobian models expressed at varying coordinate systems. Our architecture can easily be used to compute any of these Jacobians. However, the Jacobian in base coordinates is standard in the literature, thus is a good reference case for our architecture, and will be adopted in the following. For the purposes of demonstration, we use the example of an eight revolute joint arm. However, our architecture is not limited in the number or type of joints.

4.1 Rotations using CORDIC arithmetic

The CORDIC algorithms provide a fast hardware method to calculate vector rotations and inverse tangents. These algorithms were initially described by Volder [18] and further extended to hyperbolic functions by Walther.

The CORDIC iteration equations which are to be implemented in hardware are:

$$\begin{aligned} x_{i+1} &= x_i + \delta_i y_i 2^{-i}, \\ y_{i+1} &= y_i - \delta_i x_i 2^{-i}, \\ z_{i+1} &= z_i + \delta_i \theta_i, \end{aligned} \quad (5)$$

The variable, z_i , contains the total rotation angle applied, θ_i is the current rotation angle increment, and $\delta_i = \pm 1$. Through the appropriate selection of each δ_i , either the initial z_0 value can be reduced to zero (vector rotation) or the initial y_0 value can be reduced to zero (inverse tangent). These equations can be implemented with simple structures: registers, shifters, adders, and a small ROM. Figure 2 presents a block diagram of the CORDIC processor. The CORDIC algorithms have been applied by Cavallaro and Luk [2] to the Jacobi method for the SVD to produce a processor architecture which is currently being implemented by us [11] at Rice University. For the current eight joint robot architecture, we require an array to compute an 8×8 SVD. This can be performed on a 4×4 array

Figure 2: CORDIC arithmetic processor elements

of CORDIC SVD processors where each unit operates on a 2×2 submatrix. Additionally, at each iteration, prior to taking the SVD, individual CORDIC processors within the array are used (as described in the following section, see Figure 3) to compute rotations required for the DKS and Jacobian. The operation of each CORDIC unit and data flow between it and the DSP chip are shown in Figure 4.

4.2 DSP / Parallel CORDIC Jacobian

At each iteration, the values of the 8 joint variables θ are streamed in parallel to generate the following n matrices $[R_{i-1}^i]$ and n vectors $\underline{d}_{i-1}^i$:

$$[R_{i-1}^i] = \begin{bmatrix} \cos \theta_i & \sin \theta_i & 0 \\ -\sin \theta_i & \cos \theta_i & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \alpha_i & \sin \alpha_i \\ 0 & -\sin \alpha_i & \cos \alpha_i \end{bmatrix} \quad (6)$$

$$\underline{d}_{i-1}^i = \begin{bmatrix} a_i \cos \theta_i \\ a_i \sin \theta_i \\ d_i \end{bmatrix} \quad (7)$$

In the above, α_i , a_i , and d_i are constants related to the dimensions of the robot. These quantities must then be used to form the Jacobian $[J]$, which takes the form (assuming revolute joints):

$$\begin{bmatrix} \underline{z}_0^o \times \underline{d}_0^8 & \dots & \underline{z}_0^i \times (\underline{d}_0^8 - \underline{d}_0^i) & \dots & \underline{z}_0^{n-1} \times (\underline{d}_0^8 - \underline{d}_0^{n-1}) \\ \underline{z}_0^o & \dots & \underline{z}_0^i & \dots & \underline{z}_0^{n-1} \end{bmatrix} \quad (8)$$

where $\times$ denotes the cross product operation. The elements in the Jacobian are formed recursively from the above quantities $\underline{d}_{i-1}^i$ and $[R_{i-1}^i]$ as follows:

$$[R_o^i] = [R_o^{i-1}][R_{i-1}^i] \quad (9)$$

$$\underline{d}_o^i = \underline{d}_o^{i-1} + [R_o^{i-1}]\underline{d}_{i-1}^i \quad (10)$$

for $i = 1, \dots, 8$ where the initial conditions are $\underline{d}_o^0 = \underline{0}$ and $[R_o^0] = [I]$ with $[I]$ the 3×3 identity matrix and the $\underline{z}_0^i$ vectors are the third columns of the $[R_o^i]$ matrices.

A CORDIC array to compute the Jacobian is shown in Figure 3. This logical array is formed using physical sub-units of our special-purpose CORDIC SVD array. From (8), the top and bottom halves of J are

Figure 3: Dataflow for Jacobian matrix formation using CORDIC rotators.

formed differently. Both require the computation of (9) in the double triangular array shown in Figure 3. The d vectors required for the upper portion of the Jacobian are streamed in the lower triangular array. The physical array is composed of nine CORDIC rotation modules operating in parallel, with the initial seeds shown at the bottom of the array. The loading and unloading/shifting of data is performed by the DSP Processor. A diagram of the rotations produced by each module, with the data exchanges performed by the DSP processor, is shown in Figure 4. The total propagation time for this task will be 16 CORDIC rotation cycles using nine CORDIC processors in parallel. CORDIC scale factor correction will also be performed as the rotations are being applied. Notice that we have full utilization of the array.

When we seed the upper triangular array with $[0, 0, 1]^T$ then we generate the z 's. Similarly, equation (10) is seeded in the lower triangular array by d 's in (7) (notice we need only seed by $[a_i, 0, 0]^T$ to produce the d_{i-1}^i 's in one rotation, with a sign change in the DSP). Notice that by performing a succession of 2×2 rotations and holding the third vector element in register, we make a computational savings in the matrix rotations (we do not explicitly perform the two multiplications by 1 in (6)).

The subtractions and cross products to form the final Jacobian are formed by the DSP processor. The appropriate columns are then collected to form the Jacobian matrix. Other methods have been proposed to compute the Jacobian [10, 14, 16, 22]. Our approach adds to this body of techniques, using CORDIC, and

Figure 4: CORDIC co-processor organization for Jacobian calculation.

making full use of the rotations inherent in the algorithm.

In parallel with the formulation of the Jacobian, the remainder of the direct kinematics solution (x_0^8 and y_0^8) is performed by appropriate rotations of $[1, 0, 0]^T$ and $[0, 1, 0]^T$ in two other CORDIC units (shown on the left in Figure 3). When coupled with z_0^8 and d_0^8 from the Jacobian calculation, the entire direct kinematics is formed in parallel with the Jacobian calculations (recall that

$$[H_0^8] = \begin{bmatrix} x_0^8 & y_0^8 & z_0^8 & d_0^8 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (11)$$

in (2)). Notice that the total direct kinematics solution is performed using four CORDIC units, with the Jacobian calculations occupying nine units. The decomposition depicted above is not the only possible method. For example, we could perform the Jacobian operation using eight modules in parallel (shifting the triangular portions). However, in this case, the operation would take nine steps in Figure 3 instead of eight. Previous architectures to compute the DKS [10, 13] did not use CORDIC to compute the rotations. By using the CORDIC units, we are making full use of our special-purpose CORDIC SVD array.

5 Jacobian Pseudoinverse Calculation

In general, the most efficient means of computing a pseudoinverse is by its Singular Value Decomposition. Previous work in computing the pseudoinverse did not use the SVD or CORDIC [3, 15].

Figure 5: Brent Luk Van Loan array for the SVD.

5.1 SVD using CORDIC

The Singular Value Decomposition (SVD) represents an appealing method of computing the pseudoinverse. The SVD is more closely mapped onto hardware through the use of the CORDIC algorithms. In this paper, we build upon our previous work on the application of CORDIC arithmetic to the SVD, by taking advantage of the features offered by the CORDIC SVD in efficiently mapping algorithms onto hardware.

The singular value decomposition [6] of a $p \times p$ matrix M is

$$M = U \Sigma V^T, \quad (12)$$

where U and V are orthogonal matrices and Σ is a diagonal matrix of singular values. The SVD is a computationally complex algorithm which benefits from a VLSI parallel array.

The systolic array of Brent, Luk, and Van Loan [1] uses a square array of processors to implement the parallel Jacobi Method for the SVD. Figure 5 shows a typical array architecture to perform a 6×6 SVD. In the Brent-Luk-Van Loan array, the matrix is divided into 2×2 submatrices. Each processor element contains a 2×2 submatrix. The array architecture is scalable. There are two types of data flowing in this array. Rotation angles generated by the diagonal processors flow systolically along the rows and columns of the array. Matrix data elements are exchanged diagonally, after the diagonal neighbor has received and applied the necessary rotation angles. This leads to "waves" of activity moving diagonally away from the main array diagonal.

A 2×2 SVD can be described as

$$R(\theta_l)^T \begin{bmatrix} a & b \\ c & d \end{bmatrix} R(\theta_r) = \begin{bmatrix} \psi_1 & 0 \\ 0 & \psi_2 \end{bmatrix}, \quad (13)$$

where θ_l and θ_r are the left and right rotation angles, respectively. The rotation matrix is

$$R(\theta) = \begin{bmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{bmatrix}, \quad (14)$$

and the input matrix is

$$M = \begin{bmatrix} a & b \\ c & d \end{bmatrix}. \quad (15)$$

In a typical computer algorithm for the SVD, the sines and cosines of the rotation angles are computed through formulas that require division and square root operations. The explicit angles are not required and only the sines and cosines are computed. The rotations are then applied to the 2×2 submatrix using standard matrix multiplication techniques. However, time-consuming operations such as multiplication, division, and square root are needed.

As mentioned above, the CORDIC algorithms provide a fast hardware method to calculate vector rotations and inverse tangents which are essential operations for the SVD. We integrate the CORDIC SVD processor array with a DSP processor to compute inverse kinematics for redundant robots, a more general problem going beyond the earlier work on CORDIC for non-redundant robots [9, 12].

The pseudoinverse is easily expressed in terms of the information resulting from the SVD as follows:

$$J^+ = V \Sigma^+ U^T \quad (16)$$

where U and V are defined from the SVD, and Σ^+ is an $n \times m$ matrix with reciprocals of the singular values in Σ along the leading diagonal, and zeros elsewhere. The singular values of the Jacobian are obtained from the CORDIC SVD array. The current CORDIC SVD processor can be enhanced to collect the matrices U and V . The reciprocals of the singular values required for Σ^+ are generated from Σ by the DSP chip. Finally, the product $J^+ = V \Sigma^+ U^T$ is formed in the DSP chip. Part of our ongoing research concerns calculation of different schemes for selecting $\underline{\epsilon}_i$ within the DSP processor. This is possible due to the ability of DSP chips to handle differing algorithms (notice that most schemes for selecting $\underline{\epsilon}_i$ require values of $\underline{\theta}_{i-1}$ and/or $\underline{\theta}_i$, which are available locally in our architecture). This allows the entire inverse kinematics solution to be contained within the Inverse Kinematics Engine, and separate from the host machine.

After the completion of the SVD, several steps which involve a series of matrix-vector multiplications occur in (4). These are implemented using the DSP processor. The top portion of Figure 6 shows the nec-

Figure 6: CORDIC Inverse Kinematics Engine (IKE) Algorithm.

essary operations. The matrix-vector multiplications necessary to generate the Null Space projector will be performed by the DSP processor along with the Pseudoinverse computation. The DSP processor also forms the integration required to produce the joint angles from the joint velocities. The overall algorithm is configured as follows:

```

Algorithm Parallel Inverse Kinematics()
begin;
  compute Jacobian;
  compute SVD;
  do (Pseudoinverse, Null Space Projector);
  do ( $[J^+] \dot{x}$ ,  $[I - J^+ J] \epsilon$ );
  do ( $\dot{\theta} = [J^+] \dot{x} + [I - J^+ J] \epsilon$ );
  output  $\dot{\theta}$ ;
end;

```

The programmability of the DSP processor allows the use of different integration routines, such as the Adams predictor-corrector [10], Runge-Kutta methods, etc., to accuracy as desired.

Figure 7: Utilization of Processor Elements of CORDIC SVD array.

6 VLSI Implementation Issues

In our architecture, the CORDIC co-processor array is utilized for a number of functions as shown in Figure 7. Although composed of special-purpose chips, each CORDIC SVD chip (CSVD) can act as a simple CORDIC co-processor unit or as part of the systolic array for the SVD. Under the control of the DSP chip, the 9 CSVD chips in the lower righthand corner of Figure 7 compute parameters for the Jacobian matrix. Two of these nine plus two additional CSVD chips in parallel compute the DKS. At a later point in the inverse kinematics algorithm, all 16 CSVD chips compute the SVD of the Jacobian.

6.1 CORDIC Array Processor Element

The CORDIC SVD processor is implemented in a $2\mu\text{m}$ CMOS technology and fabricated through the NSF/DARPA MOSIS service. The chip measures $6900 \times 5600\mu\text{m}$ and contains over 26000 transistors [11]. The chip is primarily designed to implement the 2×2 SVD systolic algorithm, but performs in several different modes in response to a control word. The chip is also capable of acting as a dedicated CORDIC co-processor and in fact contains two 16-bit fixed-point CORDIC units. This first generation CSVD chip uses a conservative two-phase non-overlapping clocking strategy and is fully functional at 2MHz . The design uses custom adder and barrel shifter cells designed with the *magic* layout editor. The control units are implemented as PLAs and the higher level routing was completed with *octtools*.

Table 1: Performance of Inverse Kinematics Algorithm

Major Subroutines	AT&T DSP32 Ref. [15] Serial	TMS320C40 Initial 'C' Serial	AT&T DSP32 Ref. [15] 4 Parallel	TMS320C40 Assembly 1st. Gen. CSVD Array	TMS320C40 Assembly 2nd. Gen. CSVD Array
SVD	6.2ms	1.3502ms	1.7ms	1.28ms	256 μ s
Pseudo-Inverse	0.5ms	0.3234ms	0.2ms	135 μ s	135 μ s
Total	6.7ms	1.6736ms	1.9ms	1.45ms	400 μ s

We are currently designing a second generation chip which will reduce the delays in the 16-bit barrel shifter section of the CORDIC units. We expect to improve the performance of the chip so that it would operate at 10MHz. The CSVD chip communication ports are also being modified to interconnect with the TMS320C40 8 bit communication ports.

6.2 Digital Signal Processors

The current generation of floating-point DSP chips combines the flexibility of a general purpose microprocessor with high performance arithmetic and communication facilities. The TMS320C30 chip has a 60ns processor cycle time and hardware support for single cycle floating-point multiplication and addition. This leads to a high speed multiply and accumulate operation important for many of the matrix products needed for inverse kinematics. We have been using a TMS320C30 evaluation module board as part of our implementation. The current generation TMS320C40 has a 40ns processor cycle time and six hardware communication ports to support multiprocessing [4] and will provide a parallel interface to the CORDIC SVD array. We are currently working with a simulator for the TMS320C40, and projected performance results of our robotic algorithms using the TMS320C40 are presented here.

6.3 Performance Estimates

An inverse kinematics algorithm for a PUMA robot augmented with a seventh joint was implemented in the 'C' programming language. The code was compiled and loaded to maximize the use of the on-chip RAM, the zero wait-state local and global memory, the on-chip instruction cache, and the register-argument function model. The program was benchmarked on the TMS320C40 simulator and the results for each of the major subroutines and the total time are presented in Table 1. The table also contains data for the AT&T DSP32 chip presented by Maciejewski [15] using an updating scheme for the SVD. Our code computes the complete SVD using a Golub-Reinsch algorithm for the serial case on the TMS320C40 and a systolic Jacobi based algorithm for the CORDIC SVD array.

The performance of the algorithm could be improved by replacing the matrix multiplication and

transpose routines with assembly language subroutines. For instance, multiplication of two 7×7 matrices requires 2500 TMS320C40 processor cycles or 100 μ s using a 'C' language algorithm. Optimized assembly language code [4] for the same problem would require approximately 672 cycles or 27 μ s. Further speedup could be achieved by using a parallel array of TMS320C40 chips.

In our proposed architecture, we plan to replace the SVD subroutine which consumes over 75% of the execution time of the 'C' language inverse kinematics algorithm with the custom VLSI CORDIC SVD array. This array requires time for loading, computation, flushing, and unloading. The SVD algorithm requires $O(\log n)$ computation sweeps for convergence. For the 8×8 SVD, five sweeps will guarantee convergence. Based on the internal cycle counts, it will take 2560 cycles to compute the SVD, including communication with the host DSP chip. The actual time can be found by multiplying the cycle count by the reciprocal of the maximum frequency. An array of 16 first generation CSVD chips which operate at 2MHz will therefore compute the SVD in 1.28ms. The second generation CSVD chip should achieve an increase in clock frequency to 10MHz and compute the SVD of the Jacobian in 256 μ s. Even higher performance is possible if the chip were to be fabricated in the recently announced MOSIS 0.8 μ m n-well CMOS process. The cost of the new process is somewhat beyond the range of academic proof of concept implementations.

As a final performance estimate, we consider an array of second generation CSVD chips and the TMS320C40 programmed in assembly language. From Figure 6, the work which remains beyond the SVD is the reciprocal of the singular values to form Σ^+ , three matrix multiplications, two matrix vector multiplications, and a vector addition. This can be conservatively approximated by the time for five matrix multiplications. Using the above estimate of 27 μ s for assembly language multiplication, yields 135 μ s. The total time for a seven joint robot will therefore be 400 μ s.

7 Conclusions

In this paper, we describe and detail a new architecture for efficiently computing the kinematic equations (direct and inverse) for kinematically redundant robots. The architecture combines a CORDIC pro-

cessor array and a programmable DSP processor in a novel and interesting manner, fusing together several systolic algorithms. We detail a new parallel method to compute both the Direct Kinematics Solution and the manipulator Jacobian in parallel using units of the CORDIC array. The array, which acts as a co-processor to the DSP chip, is then used to compute the SVD of the Jacobian and its pseudoinverse. The DSP processor is used to compute the joint velocities and positions to complete the inverse kinematics solution. The architecture is valid for nonredundant as well as redundant arms, and can use any type of high performance DSP chip. We have constructed the first generation CORDIC SVD chip and are currently building, testing, and integrating the CORDIC SVD array with the Texas Instruments TMS320C40 processor. We estimate the time for each iteration of the inverse kinematics of an eight joint arm using our architecture to be $440\mu\text{s}$. This suggests that some of the more complex schemes for redundancy resolution suggested in the literature can become tractable in real-time using our architecture.

Acknowledgments

This work was supported in part by NSF under grants MIP-8909498, DDM-9202639, and MSS-9024391 and by DOE Sandia National Laboratory Contract #18-4379A. The authors would like to thank Dr. W.A. Gordon, R.M. Piedra, G.S. Thomas and R.H. Price of the Digital Signal Processing Division of Texas Instruments, Houston.

References

- [1] R. P. Brent, F. T. Luk, and C. F. Van Loan. Computation of the Singular Value Decomposition Using Mesh-Connected Processors. *Journal of VLSI and Computer Systems*, 1(3):242-270, 1985.
- [2] J. R. Cavallaro and F. T. Luk. CORDIC Arithmetic for an SVD Processor. *Journal of Parallel and Distributed Computing*, 5(3):271-290, June 1988.
- [3] P.R. Chang and C.S.G. Lee. Residue Arithmetic VLSI Array Architecture for Manipulator Pseudo-Inverse Jacobian Computation. *IEEE Transactions on Robotics and Automation*, 5(5):569-582, 1989.
- [4] D.C. Chen and R.H. Price. A Real-Time TMS320C40 Based Parallel System for High Rate Digital Signal Processing. In *Proceedings 1991 IEEE ICASSP*, pages 1573-1576, Toronto, Canada, 1991.
- [5] B. W. Drake and T. C. Hsia. Implementation of a Unified Robot Kinematics and Inverse Dynamics Algorithm on a DSP Chip. In *Proceedings 1992 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1193-1200, Raleigh, NC, 1992.
- [6] G. H. Golub and C. F. Van Loan. *Matrix Computations, Second Edition*. Johns Hopkins Univ. Press, Baltimore, MD, 1989.
- [7] J.H. Graham. Special Computer Architectures for Robotics: Tutorial and Survey. *IEEE Transactions on Robotics and Automation*, 5(5):543-554, October 1989.
- [8] R.T. Gregory and E.V. Krishnamurthy. *Methods and Applications of Error-Free Computation*. Springer-Verlag, New York, 1984.
- [9] R.G. Harber, J. Li, X. Hu, and S.C. Bass. The Application of Bit-Serial CORDIC Computational Units to the Design of Inverse Kinematics Processors. In *Proceedings 1988 IEEE Conference on Robotics and Automation*, pages 1152-1157, Philadelphia, PA, 1988.
- [10] T.C. Hsia, K.W. Current, Z. Mao, W.S. Chu, J. Liu, G.Z. Lu, and W.H. Han. A Proposed New VLSI Architecture for Real-Time Robot Manipulator Control. *International Journal of Robotics and Automation*, 6(4):169-178, 1991.
- [11] K. Kota and J.R. Cavallaro. Numerical Accuracy and Hardware Tradeoffs for CORDIC Arithmetic for Special-Purpose Processors. *IEEE Transactions on Computers*, 1993. To appear.
- [12] C.S.G. Lee and P.R. Chang. A Maximum Pipelined CORDIC Architecture for Inverse Kinematic Position Computation. *IEEE Journal of Robotics and Automation*, RA-3(5):445-458, October 1987.
- [13] S.S. Leung and M.A. Shanblatt. Real-Time DKS on a Single Chip. *IEEE Transactions on Robotics and Automation*, 3(4):281-290, 1987.
- [14] C.J. Li, A. Hemami, and T.S. Sankar. An Efficient Computational Method of the Jacobian for Robot Manipulators. *Robotica*, 9:231-234, 1991.
- [15] A.A. Maciejewski and J.M. Reagin. A Parallel Algorithm and Architecture for the Control of Kinetically Redundant Manipulators. In *Proceedings 1992 IEEE Conference on Robotics and Automation*, pages 488-493, Nice, France, 1992.
- [16] D.E. Orin and W.W. Schrader. Efficient Computation of the Jacobian for Robot Manipulators. *International Journal of Robotics Research*, 3(4):66-75, 1984.
- [17] B. Siciliano. Kinematic Control of Redundant Robot Manipulators: A Tutorial. *Journal of Intelligent and Robotic Systems*, 3(3):201-210, 1990.
- [18] J. Volder. The CORDIC Trigonometric Computing Technique. *IRE Trans. Electronic Computers*, EC-8(3):330-334, Sept. 1959.
- [19] I.D. Walker and J.R. Cavallaro. Parallel VLSI Architectures for Real-Time Control of Redundant Robots. In *Proceedings 1991 American Nuclear Society Meeting on Robotics and Remote Systems*, pages 299-310, Albuquerque, NM, 1991.
- [20] I.D. Walker and J.R. Cavallaro. Parallel VLSI Architectures for Real-Time Kinematics of Redundant Robots. *Journal of Intelligent and Robotic Systems*, 1993. To appear.
- [21] Y. Wang and S.E. Butner. A New Architecture for Robot Control. In *Proceedings 1987 IEEE Conference on Robotics and Automation*, pages 664-670, Raleigh, NC, 1987.
- [22] T.B. Yeung and C.S.G. Lee. Efficient Parallel Algorithms and VLSI Architectures for Manipulator Jacobian Calculation. *IEEE Transactions on Systems, Man, and Cybernetics*, 19(5):1154-1166, 1989.