

A MASSIVELY PARALLEL ACCELERATOR FOR THE SAS® SYSTEM AT THE CENSUS BUREAU

Mr. Jon Michael Fortier
U. S. Bureau of the Census

Mr. Douglas Burns
MasPar Computer Corporation

Abstract

The ability to summarize large volumes of data is limited by traditional computer systems which can take hours of CPU time and days of elapsed time to process multi-gigabyte datasets. The Census Bureau has begun testing of an alternative approach using a massively parallel computer from MasPar Computer Corporation which is accessed via a SAS/AF® application called TABLE-MAKER®. Initial results indicate that CPU and I/O-intensive analyses of the census data take a small fraction of the time required by traditional computer systems when performed on the massively parallel computer.

Introduction

This paper describes preliminary results obtained from tests currently underway at the University of Maryland Institute for Advanced Computer Studies (UMIACS). The tests are being performed in order to characterize the performance of a massively parallel computer in delivering summary statistics from large Census datasets. Two three-way tabulations were performed.

The testing is a cooperative effort between MasPar Computer Corporation, UMIACS, the University of Michigan Population Study Center, and other researchers using Census data. The final results of these tests will be reported at SUGI 18 in May '93.

Description of Data

Two datasets are used in this analysis; the Current Population Study (CPS) from March '89, '90 and '91, and the Public Use Microdata Sample (PUMS) 5% sample from 1980. Both datasets are collected and maintained by the U. S. Bureau of the Census.

CPS Data

The CPS is a monthly survey of the population and is comprised of a consistent core survey and a monthly supplemental survey where each monthly survey has a different focus. The March survey supplemental topic is "income" which tends to be the largest monthly survey. The data file has a tri-level hierarchy of 'household', 'family' and 'person' records.

The composition of the CPS dataset is shown in Table 1 below.

Table 1

CPS Dataset Breakdown
(number of records)

	1989	1990	1991
Household	70,454	75,269	75,076
Family	60,629	66,228	66,459
Person	144,687	158,079	158,477
Total	275,770	299,576	300,012
Megabytes	181	197	197

Each CPS record is 658 bytes long, including CR/LF.

PUMS Data

PUMS is the largest publicly available sample of US Census data. The PUMS dataset used in this analysis is the 5% sample of the 1980 decennial census. It consists of approximately 10.8 million records containing 244 fields per record. The PUMS database, which is pre-sorted by State, employs a two-level hierarchical data structure, with 'household' and 'person' records. The entire dataset consumes approximately four gigabytes of disk space.

Description of Computing Task

Data Preparation

In order to make most effective use of the massively parallel computer, the database is loaded onto the high-speed parallel disk array, and the data structure is altered. Preparation of the CPS and PUMS data to allow access by the SAS/AF® application, TM/ACCESS® (see "Description of the Software"), involved two steps:

- 1) flattening of the hierarchical records into a flat file
- 2) writing each field into a separate binary file

Simple tools included with TABLE-MAKER® (see "Description of the Software") automated this one-time flattening and inversion. The data conversion and loading was performed on the MasPar system overnight on the entire four gigabytes. This resulted in some growth of the data storage requirement, the growth being roughly proportional to the number of levels of hierarchy in the data.

Cross Tabulation

The computational tasks were simple cross-tabulations of a small number of fields from the datasets. The point-and-click interface of TM/ACCESS provided a simple means of defining the tables to be generated by allowing the user to select datasets and variables of interest with a very comfortable and intuitive graphical interface. A TABLE-MAKER command file was generated by TM/ACCESS that drove the tabulation and extraction of data on the MasPar system.

A three-way tabulation of AGE by SEX by RACE was performed on the 1980 5% PUMS dataset of 10,800,000 records. The table shows (see Table 2) the counts of males and females in ten age-groups by four race classifications. The results of the tabulation were returned to SAS for further manipulation and reporting.

Table 2: Population Counts from PUMS

Age (years)	Sex	Race			
		White	Black	Hispanic	Other
0 to 9	Male	667,185	128,113	29,978	30,131
	Female	633,406	125,310	29,066	29,540
10 to 19	Male	806,490	143,802	27,909	28,648
	Female	773,057	144,668	26,965	27,559
20 to 29	Male	828,533	116,892	29,665	29,741
	Female	829,912	132,975	27,423	31,110
30 to 39	Male	659,353	76,362	17,080	24,556
	Female	668,162	91,821	17,037	26,892
40 to 49	Male	472,954	53,529	9,924	14,941
	Female	489,741	65,662	10,564	16,196
50 to 59	Male	489,253	47,827	6,623	10,350
	Female	533,688	60,036	7,317	12,045
60 to 69	Male	383,963	35,644	3,049	6,526
	Female	459,712	46,388	3,727	7,502
70 to 79	Male	210,580	18,894	1,590	3,683
	Female	310,998	28,393	2,016	4,106
80 to 89	Male	67,462	5,343	417	986
	Female	135,593	9,331	549	1,541
90 to 99	Male	8,955	973	75	197
	Female	22,626	2,003	98	209

A three-way tabulation of employment status ("WRKSTATUS") by SEX by STATE-GROUP was performed on the CPS dataset. The table shows (see Table 3) the counts of males and females in three employment classifications (employed, unemployed, not in workforce) by seven state-groups.

Table 3: Employment Status from CPS

State Group	Sex	Work Status		
		Employed	Unemployed	not in Workforce
1	Male	3,350	260	2,496
	Female	3,067	174	3,329
2	Male	6,095	368	5,194
	Female	5,110	250	7,404
3	Male	5,843	417	4,911
	Female	4,970	312	6,730
4	Male	3,431	208	2,754
	Female	3,028	128	3,638
5	Male	6,705	311	5,522
	Female	6,038	291	7,605
6	Male	1,685	102	1,655
	Female	1,529	93	2,307
7	Male	12,591	827	11,406
	Female	10,380	665	14,900
other	Male	39,700	2,493	33,938
	Female	34,122	1,913	45,913

Description of the Hardware

The massively parallel computer system from MasPar consists of an array of from 1,024 to 16,384 processors which work together to provide 1,300 to 68,000 MIPS for compute- and data-intensive problems (see Figure 1). Main memory ranges from 16 megabytes to 1 gigabyte with up to one gigabyte of I/O RAM or 'RAM disk'. The massively parallel computer utilizes an industry-standard UNIX® workstation as a front-end processor. The system used in the tests described in this paper has 4,096 processors and 256 megabytes of memory.




Figure 1

The parallel disk array is a NFS-mountable, RAID3 configuration capable of reading and writing single files of up to 2 gigabytes through high-speed I/O channel connections to the parallel processor. The capacity of a single array ranges from 5.5 gigabytes to 33 gigabytes. A single array can transfer data into the massively parallel processor at sustained rates of 30 megabytes per second. Disk array size can be scaled to over 500 gigabytes with transfer rates of up to 200 megabytes per second for a single parallel processor. A single array of 22 gigabytes was used in these tests.

Description of the Software

SAS runs on the UNIX subsystem (see Figure 1) and provides the interface and tools to access and use the massively parallel computer. A SAS/AF application called TM/Access is opened from within SAS to run TABLE-MAKER, the software application that drives the MasPar computer. TM/ACCESS is written in SAS Language® and SAS SCL®. TABLE-MAKER is written in MPL®, MasPar's parallel version of ANSI C.

TABLE-MAKER may be accessed locally or remotely by TM/ACCESS. TABLE-MAKER, a product of Chisq Ltd. of the United Kingdom, provides two major classes of commands; "Tabulate" and "Extract". "Tabulate" provides summary statistics including frequency counts, sums, sums of squares, min, max, and median. "Extract" provides the summary results to SAS for further processing and reporting. The syntax is relatively simple to learn and operation is facilitated by the point-and-click interface provided by TM/ACCESS.

The existence and operation of the MasPar system was completely transparent. From a user's perspective, there was no requirement to learn anything beyond the operation of the SAS/AF application, TM/ACCESS. However, the following short description of the parallel programming concepts underlying TABLE-MAKER provides a degree of insight into the operation of this powerful system.

The design and programming methodologies applied to massively parallel computers are significantly different than those used on serial and modestly parallel systems. The parallel programming model employed by the MasPar system assumes that there is a large dataset to be processed and that there is a single processor for each data element in the set. The same set of instructions is applied to all the processors. For example, in a simple MIN or MAX operation, the MasPar system can compare up to 16,384 values simultaneously.

Logical conditions in each processor allow for selective computation by each processor on its own local data. If data held by a given processor does not match the criteria for execution of the instruction being applied to the array, that processor is disabled and does not execute that instruction. For example, the frequency counts of individuals within a certain range of ages can be accomplished even though some of the processors contain data that does not satisfy the criteria.

Preliminary Results of Analysis

Results of the tabulation of "AGE" by "SEX" by "RACE" on the 10.8 million-record, 3.5 gigabyte PUMS dataset were returned in 55 seconds. This same tabulation takes nearly 30 minutes on the most powerful mainframes supported by SAS.

Results of the tabulation of "WRKSTATUS" by "SEX" by "AGE" on the 600 megabyte, 875,000-record CPS dataset were returned in less than one second. Such a cross-tabulation performed on the uni-processor systems currently in use at the Census Bureau take from 30 to 100 seconds to complete.

Conclusions

It has been demonstrated through these tests that SAS, through a SAS/AF application, has the ability to rapidly process large volumes of Census data. Ultimately, the integration and availability of massively parallel systems to the Census Bureau's networked user population may result in a substantial increase in the amount of analysis performed.

The combination of SAS, TABLE-MAKER, and the MasPar system applied to the analysis of Census data has shown two things:

1) that the analysis of large datasets can be accelerated substantially over currently used methods and platforms, and

2) that this powerful level of performance is completely and transparently accessible from a familiar and widely-used software system, SAS.

Massively parallel computers such as the MasPar system used in this study offer significant benefits to analysts working with large volumes of data. Quicker turn-around times and increased accessibility are the primary benefits. However, the broad adoption of a massively parallel computing platform depends on the ease-of-use of these systems and, therefore, requires the robust facilities of a widely used and familiar application software environment such as SAS.

We believe the rapid data-analysis capability of a massively parallel computer combined with the robust functionality of SAS will usher in a new paradigm for the analysis of large datasets. This paradigm is defined by a level of interactivity and flexibility that will encourage a much greater level of exploratory data-analysis and, ultimately, a much greater understanding of the information contained in that data.

SAS, SAS/AF, SAS Language and SAS SCL are registered trademarks of The SAS Institute.

TABLE-MAKER and TM/ACCESS are registered trademarks of Chisq Ltd.

MasPar, MP-1 and MPL are registered trademarks of MasPar Computer.