VI. Parallel Computing: a Latin American Perspective

S. W. Song

Universidade de São Paulo
Instituto de Matemática e Estatística
Departamento de Ciência da Computação
CEP 05508-900 São Paulo, SP, Brazil
e-mail: song@ime.usp.br

1 Introduction

High-performance computing has almost always been associated with the most expressive new results obtained in various areas of Science and Technology. The appearance of the so-called supercomputers has opened new horizons for scientific research and development, enabling observations and experimentations considered until then impossible or economically infeasible.

High-performance computing is intimately related to the solution of the so-called "Grand Challenges" that include such problems as Forecasting of Weather and Climate, Material Sciences, Structural Biology, Chromodynamics, Transport, Ocean Sciences, Medicine and Health.

According to the recent (November 10, 1995) TOP500 Supercomputer Sites report [12], the topmost 500 supercomputers are distributed as follows.

Table 1

Continent or country	Number of supercomputers in the TOP500 list	
Unites States	269	54%
Europe	139	28%
Japan	73	14%
Others	19	4%
Total	500	100%

Notice the relatively small number (4%) of supercomputer sites located outside the U.S., Europe and Japan. Among these 19 supercomputer sites, only two are located in Latin America, one in Mexico and one in Brazil. At the Universidad Autónoma Metropolitana, Istapalapa, Mexico, the Silicon Graphics Power Challenge occupies the 268th position among the TOP500 list. In Brazil's INPE (Instituto Nacional de Pesquisas Espaciais – National Institute of Space Research), the Nippon Electric Company SX-3/12R occupies the 420th position. This machine is used in weather forecasting research. The following table gives a summary.

Table 2

(Source: TOP500 Supercomputer Sites – November 10, 1995 report [12].)

Country	Institution	Equipment	Usage	Ranking
Mexico	Univ. Aut. Metropolitana	SGI Power Challenge	Academic	268
Brazil	INPE / CPTEC	NEC SX-3/12R	Weather	420

The number of supercomputer sites located in Latin America has been a little larger two years ago. According to the 1993 TOP500 Supercomputer Sites report, there were two supercomputers located in Brazil e another two in Mexico, as shown in the following table.

Table 3

(Source: TOP500 Supercomputer Sites – November 11, 1993 report [11].)

Country	Institution	Equipment	Usage	Ranking
Brazil	INPE / CPTEC	NEC SX-3/12R	Weather	93
Mexico	Univ. Nac. Aut. de México	Cray Y-MP1/132	Academic	292
Mexico	ITESM	IBM 9076-001 SP-1	Academic	408
Brazil	UFRGS	Cray Y-MP2E/232	Academic	461

Evolution of the TOP500 equipments in terms of architecture is of special interest. Notice the gradual and steady increase of the MPP (massively parallel processor) technology during the past three years (see the following table [12]). It is quite certain that Parallel Computing should play a very important role in supercomputing in the next decade.

Table 4

Month/year	Number of MPP systems in the TOP500 list	
Jun/1993	156	31%
Nov/1993	187	37%
Jun/1994	227	45%
Nov/1994	239	48%
Jun/1995	230	46%
Nov/1995	284	57%

With respect to CPU technology, off-the-shelf CMOS seems to gain ground at an extremely fast pace. The following is based on the TOP500 supercomputer sites during the past three years.

Table 5

Month/year	Number of systems using off-the-shelf CMOS	
Jun/1993	109	22%
Nov/1993	124	25%
Jun/1994	193	39%
Nov/1994	242	48%
Jun/1995	322	64%
Nov/1995	364	73%

2 Supercomputing systems in Latin America

2.1 National Supercomputer Centers

As shown in the last section, the number os supercomputers among the TOP500 list located in Latin America is very small. However there are a considerable number of supercomputers outside the TOP500 list. A very important role in this dissemination of high-performance computing is played by the federal government. In Brazil the Ministry of Science and Technology created the SINAPAD program (Sistema de Centros Nacionais de Processamento de Alto Desempenho – National High-Performance Processing Centers). This is similar to the National Science Foundation Supercomputer Centers in the United States, with, however, a much smaller budget. The mission of SINAPAD is to provide modern computing services of quality and high capacity to researchers and professionals, as a means to diffuse high-performance computing to the various segments of Science and Technology. Five such national supercomputer centers (called CENAPAD – Centro Nacional de Processamento de Alto Desempenho) have been installed or planned, in the states of Rio Grande do Sul, São Paulo, Rio de Janeiro, Minas Gerais and Ceará. One characteristics is that these centers, once installed, are supposed to be self supported through the services they provide. The central coordination of the CENAPAD is established by the Ministry of Science and Technology, through its FINEP (Financiadora de Estudos e Projetos – Research and Project Funding Agency).

The SINAPAD program is not limited to the national supercomputer centers which aim to diffuse the usage of high-performance technology. To a lesser extent, there is also a program to finance the research and design of high-performance hardware. A much smaller budget is dedicated to this end. The research and development groups deem to be adequate to this high-performance hardware design program are: the Universidade Federal do Rio de Janeiro, through its COPPE (Coordenação dos Programas de Pós-Graduação de Engenharia) and the Universidade de São Paulo, through its LSI (Laboratório de Sistemas Integráveis).

In the following we give a brief description of three of the five CENAPADs.

CESUP - RS

CESUP - RS (Centro Nacional de Supercomputação na Região Sul) was the first national supercomputer center to be inaugurated (June 1992). It is installed at the Universidade Federal do Rio Grande do Sul.

The hardware configuration of CESUP - RS is as follows.

Cray Research Y-MP 2E/232 with 2 central units, each with:

- peak performance of 330 MFLOPS
- memory of 256 Mbytes
- disk space of 16 Gbytes

CENAPAD - SP

CENAPAD - SP (Centro Nacional de Processamento de Alto Desempenho de São Paulo), the second national center (inaugurated in March 1994) is located at the Universidade Estadual de Campinas. Its hardware configurations is as follows.

IBM 9076 SP1 - 8 processors each with

- memory of 256 Mbytes
- disk space of 2 Gbytes

IBM ES 9021 - model 711

- peak performance of 563 MFLOPS
- memory of 512 Mbytes
- extended memory of 512 Mbytes

CENAPAD - RJ

CENAPAD - RJ (Centro Nacional de Processamento de Alto Desempenho do Rio de Janeiro) is installed at LNCC (Laboratório Nacional de Computação Científica). The hardware configuration is as follows.

IBM SP-2 with 16 processors each with:

- Memory of 256 Mbytes
- local disk of 2 Gbytes
- 2 data servers RS/6000 mod. 980 with 55 Gbytes
- total peak performance of 2 GFLOPS.

A planned upgrade (Feb. 96) will elevate the processing capacity to 8.7 GFLOPS.

2.2 Other supercomputer systems

A partial list of other supercomputer systems is included in Appendix A (page 16).

3 Computing in Latin America

As in the case of many third world countries, exports of products are essential in the national economy. Some of the Latin American countries have shown success in their capacity to develop and export software. As reported in [3], exports of software developed in Chile amounted to approximately 22 million dollars in 1993. In addition to application packages and software utilities, the exported products include applications dealing with mining and forestry, in which Chile has special experience.

In 1992, the Brazilian Ministry of Science and Technology, through CNPq (Conselho Nacional de Desenvolvimento Científico e Tecnológico), launched an aggressive nationwide program, called SOFTEX 2000. What follows is based on [20] which can be consulted for more details. The SOFTEX 2000 program is supported by the United Nations Development Programme (UNDP). The main goal is to redirect and steer the Brazilian industry toward the development of software of quality for export. Brazil presents an active and high-potential economy, with the 8th Gross National Product

(GNP) of the world. Many factors show the strong vocation for Informatics: this country hosts a large number of national and international software companies, with professionals trained through many high-quality university courses in Computer Science, a nationwide computer network (RNP – Rede Nacional de Pesquisa) interconnecting virtually all the major universities and research institutes. The expenditure in Informatics in Brazil, as compared to the GNP, present similar indices as observed in more developed economies as Japan, Germany and France. One notable area in the usage of Informatics and networking is in transactional banking systems, certainly among the most advanced in the world. A research by Andersen Consulting shows that one in each three homes of classes A and B has a computer. The total of computer equipments is 1.24 million, about 42% of which with fax/modem. The following table contains some interesting comparative data (Sources: Secretaria de Informática, Brazil, also [13].)

Table 6

(In billions of US dollars.)

Country	Hardware	Software products	Services in Informatics	Total	Total/GNP	Total per capita
USA	85,6	31,8	58,8	176,2	2,83	646
Japan	35,5	6,4	30,2	72,1	2,04	295
Germany	17,0	5,8	15,0	37,8	1,94	465
France	10,4	4,1	12,0	26,5	2,00	466
UK	11,7	4,2	10,2	26,1	2,41	438
Italy	6,2	3,3	7,1	16,6	1,35	292
Brazil	5,2	0,9	2,7	8,9	2,00	63
Netherlands	3,4	1,7	2,8	7,9	2,36	499

Since the end of the Informatics Reserve Market Policy in 1990, considerable amount of effort has been expended towards the commercial integration to international markets. (Refer to [17, 19] for some literature on the Brazilian national policy on Informatics.) The SOFTEX 2000 program, a joint initiative between the government and the private industry, aims to foster the export of high quality software as an strategic economic alternative by the end of this century. It is viewed as a high priority program by the government to promote software in the international market. An outpost office has already been established in Florida, U.S.A., as a reference point of the program, to facilitate access and contact with the North American market. More specifically, the goals of SOFTEX 2000 are to conquer one percent of the international software market by the end of this century, to capacitate more than one thousand companies and generate 50 thousand new qualified jobs. This ambitious program is coordinated by a special team from CNPq (Conselho Nacional de Desenvolvimento Científico e Tecnológico). It is the role of this central coordination to formulate the strategy and allocate resources of

CNPq and United Nations (UNDP). In addition to the central coordination, there are regional entities to promote the software segment in their respective regions of influence.

Among the small Latin American countries (defined as having less than ten million inhabitants), some have shown success in serving as hosts for foreign investment and as export platforms for international corporations. These corporations take advantage of cheap labor to set up data-processing centers in the Dominican Republic, Jamaica and Costa Rica [7]. The latter is host to several major international corporations to be a regional leader in furnishing international networking to Central America.

3.1 Research in Parallel Computing

A partial list of universities in which there are research on Parallel Computing follows. The mentioned areas are for illustrative purpose.

Table 7

University	Examples of Research Areas
Univ. Federal de Pernambuco, Brazil	Parallel Architectures
	Parallel Algorithms
Univ. de Brasília, Brazil	Parallel Architectures
	Concurrent programming
	Distr. Operating Systems
Univ. Fed. Mato Grosso Sul, Brazil	Parallel Algorithms
Univ. Fed. Minas Gerais, Brazil	Parallel Processing
	Parallel Algorithms
Univ. de São Paulo (Eng.), Brazil	Parallel Computer Architectures
	Parallel languages and compilers
Univ. de São Paulo (Comp. Sci.), Brazil	Parallel Algorithms
	Systolic Algorithms
Univ. Estadual Campinas, Brazil	Parallel Algorithms
	Efficient implementation issues
Univ. Fed. São Carlos, Brazil	Parallel Algorithms
Pont. Univ. Católica Rio de Janeiro	Parallel Algorithms
Unvi. Fed. Rio de Janeiro, Brazil	Parallel Computer Architectures
	Parallel Algorithms
	Superscaler Architectures
Univ. Fed. Rio Grande Sul, Brazil	Parallel Architectures
Centro de Calculo, Fac. de Ingenieria, Uruguay	Parallel Proc. on workstations
Inst. Tec. Autónomo de México, Mexico	Distributed Systems
Universidad de Chile, Chile	Parallel Algorithms
	Cellular Automata
Univ. Simón Bolívar, Venezuela	Parallel Algorithms
	Parallel Programming Tools

In Brazil the following universities have official graduate programs (recognized by the Ministry of Education) in Computer Science.

> Pontifície Universidade Católica do Rio de Janeiro, Universidade de São Paulo (campus São Paulo and São Carlos), Universidade Estadual de Campinas, Universidade Federal do Rio de Janeiro, Universidade Federal de Pernambuco, Universidade Federal de Minas Gerais, Universidade Federal do Rio Grande do Sul, Instituto Nacional de Pesquisas Espaciais, Universidade Federal da Paraíba, Universidade Federal do Paraná, Universidade Federal de Santa Catarina, Universidade Federal de São Carlos, Universidade de Brasília, Instituto Militar de Engenharia,

Among the above list, the first eight also have Ph.D. program and most of the graduate programs include the area of Parallel Computing.

In [14] a very careful study of the the various aspects of Computer Science research in Mexico is presented. It contains also some interesting statistics concerning the important question of human resources in Computer Science of various Latin American countries. Brazil has largest number of around 380 Ph.D.s in Computer Science (with one Ph.D. for every 433 thousand inhabitants), Chile has 70 (one Ph.D. for every 171 thousand inhabitants), Mexico has only one Ph.D. in Computer Science per million inhabitants. In Mexico the Ph.D.s in Computer Science are concentrated in one public and five private centers. Of the latter, four are departments of two private universities: Instituto Autónomo de México (ITAM) and Instituto Tecnologico de Estudios Superiores de Monterrey (ITESM). The other private institution is the Laboratorio Nacional de Informatica Avanzada (LANIA). The Sección de Computación of the Department of Electrical Engineering of the Instituto Politecnico Nacional is the only public center that contains a group of Ph.D. researcher in Computer Science.

3.2 Conferences and Journals in Computer Science

Most countries maintain their local computer societies, as the Sociedade Brasileira de Computação (SBC), Sociedad Chilena de Ciencia de la Computación (SCCC), Sociedad Argentina de Informática e Investigación Operativa (SADIO). There is also a Latin American Computer Society – CLEI – Centro Latino Americano de Estudios en Informática.

The Brazilian SBC organizes a main annual Computer Science Conference (with an attendance of more than 2 thousand in the recent August 1995 edition), as well as about a dozen specific symposia in such areas as Databases, Software Engineering, Computer Graphics and Image Processing, Artificial Intelligence, Computer Networks, Conception of Integrated Circuits, Neural Networks, Music and Computers, Informatics in Education, Programming Languages, etc. Some of these symposia have international program committees, with submissions from all around the world. Since 1987, SBC started the annual symposium on Parallel Processing, currently named SBAC-PAD – Simpósio Brasileiro de Arquitetura de Computadores e Processamento de Alto Desempenho (Computer Architecture and High-Performance Processing Symposium). The specific area of Parallel Computing has grown considerably during the past few years in Brazil. In last August 1995, the 7th SBAC-PAD was held in Canela, Rio Grande do Sul, with nearly 600 participants.

The Chilean SCCC organizes an annual conference called the International Conference of the Chilean Computer Science Society, a conference of high quality with international program committee.

In 1992, the Department of Computer Science of the Universidade de São Paulo (DCC/IME/USP) sponsored the international symposium LATIN'92 (*Latin American Theoretical INformatics*), the first of a series of symposia in Theoretical Computer Science to be held in Latin America. The conference proceedings are published as Lecture Notes in Computer Science by Springer-Verlag. The second edition of LATIN was held in 1995 in Valparaíso, Chile. The third one is planned to be held again in Brazil, at the Universidade Estadual de Campinas, in 1997.

In 1993, the First South American Workshop on String Processing was organized and held at the Universidade Federal de Minas Gerais. The second workshop was repeated in Chile, held immediately after LATIN, in April 1995. The third workshop will take place at the Universidade Federal de Pernambuco, in August 1996.

An annual Latin American conference is organized by CLEI, called Conferência Latinoamericana de Informática. This conference is held each year at a different Latin American country.

The Sociedade Brasileira de Computação (SBC) publishes the journal *Revista Brasileira de Computação*. In 1994, SBC inaugurated a new international journal, named *Journal of the Brazilian Computer Society* (JBCS), with an international editorial board. JBCS is published three time a year. The third issue, published in July 1995, was dedicated to Parallel Computing.

4 International cooperation - a new experience

The cooperation between the university and industry has been much enphasized, particularly in the European ESPRIT program. Cooperative programs involving diverse partners from the academia and industry have been adopted as the sine qua non condition in many recent international cooperation programs.

In the area of Parallel Computing, the ITDC'94 program (Information Technologies for Developing Countries) can be considered as extremely important to the development of this area in Latin American countries. ITDC was an initiative of the Commission of the European Communities (CEC). The purpose of this initiative is to provide research infrastructure, in terms of parallel computers, to aid research groups located in developing countries. The goal is to establish scientific cooperation between research institutions in developing countries and in Europe. The CEC funding, through ITDC, includes the purchase of European parallel computers, as well as other minor costs for installation, training and travel. The call for projects was responded with great enthusiasm, with more than 250 projects submitted from research groups located all around the world. The high competition ensures the quality of the more than twenty selected projects to receive funding. In Latin America, we can mention this following partial list of institutions that were selected: ITESM of Mexico, Universidad de Venezuela, Universidad de Chile, Universidade de São Paulo (two projects – one from the Department of Computer Science and another from Electrical Engineering), Universidade Estadual de Campinas (Computer Science and Electrical Engineering) Universidade Federal do Rio de Janeiro, etc.

Another CEC cooperative program involving multiple partners is the ALFA program (America Latina Formación Academica), with the purpose of encouraging scientific cooperation between European and Latin American institutions, mainly in the formation of human resources.

In 1995 CEC launched the INCO program (International Cooperation with Third Countries - Part C) with a total funding of 208.98 million ECU through 1998. The proposal should be transnational involving at least two partners from different member states of the European Union and at least one partner from a developing country, with strong preference to projects with at least two partners from separate developing countries in the same region. In each annual call for proposals, high priority areas are identified. One such area is Applications of High-Performance Computing, Parallel Programming and Networking (HPCN) with the following objectives:

- raising the awareness of and promoting the use of HPCN in industry
- training HPCN applications developers and users in industry
- establish HPCN best practice in a user environment
- stimulating the development of the HPCN infrastructure
- undertaking HPCN research as identified by application requirements.

This four-year program will undoubtly boost the already effervescent Parallel Computing research in Latin American countries to even higher levels.

One notable point to be noticed in recent funding of research projects is the presence of collaboration with industrial partners, with preference to application-oriented research topics that would result in substantial benefits to society. To illustrate this point we mention the CEC/ITDC project being coordinated by this author (Universidade de São Paulo (IME/USP), Departamento de Ciência da Computação, Grupo de Computação Paralela e Distribuída). This project has the collaboration of GMD FIRST

Berlin and Parsytec (which provided the PowerXplorer 16/32 Parallel Computing System installed at IME/USP).

One of the goals of this project is to develop application-oriented software for numerically intensive applications. The chosen application area is Meteorological Computations and Environmental Analysis (Monitoring of Pollutants Dispersion). The research part of the project in environmental and atmospheric applications will be conducted by researchers of the Department of Applied Mathematics and Department of Atmospheric Sciences of USP. We will also seek collaboration with governmental agencies for environmental control in the City and State of São Paulo. The goal of the project is to improve the understanding of the meteorological control processes on the air pollution in São Paulo and develop operational techniques for predicting local atmospheric conditions in São Paulo. The Department of Atmospheric Sciences of USP has implemented the Regional Atmospheric Modeling System (RAMS) in 1990 and it will be used as the main atmospheric simulation tool. The smog analysis and monitoring system DYMOS, developed in Germany[18], will be used for the simulation of the air pollution in São Paulo. Comparison between the atmospheric components of both DYMOS and RAMS shall be conducted. RAMS has been undergoing significant rearrangement of the code in order to run on clusters of workstations and massive parallel machines. The parallel version of RAMS was available for the users at the end of 1994 and shall be ported to the Parsytec system. As part of the ongoing collaboration between GMD and IME/USP a fine grain message passing version of the DYMOS model shall be developed.

Numerical modeling of the local circulation in São Paulo is an important technique for understanding the dispersion processes associated with the air pollution. The basic role of the sea breeze and the complex orography on the initiation of precipitation has already been explored. However, most of the heavy pollution episodes in São Paulo occur during the dry season. Thus, atmospheric modeling will be applied to Winter type typical situations in São Paulo, considering realistic distributions of surface characteristics, such as high resolution topography (order of a kilometer), type of soil and vegetation, surface roughness, moisture and albedo. Besides the processes studies, an operational technique for forecasting the local flow shall be developed on a nested version of RAMS with coarse grid resolution of 32 km and a fine grid of 2 km covering the metropolitan area of São Paulo.

5 Design of scalable algorithms

As a joint work of this author with F. Dehne (School of Computer Science, Carleton University, Ottawa, Canada), we have investigated the problem and issues of efficient parallel implementations. The following is extracted from an ongoing joint project.

The main problem in parallel computing is the well known "software bottleneck". Current commercial applications of parallel machines are still restricted mainly to trivially parallelizable problems where communication requirements are obviously low. On the other hand, there is a large body of literature on parallel algorithm design for many non-trivial problems. However, these results suffer from the fact that there exists no agreed model of parallelism that is close enough to existing machines to allow for a reasonable prediction of the speed of an implementation. The problem is obvious for PRAM based parallel algorithms, but even network based parallel algorithms are often very problematic and the speed obtained when implementing such algorithms on a commercial multiprocessor is often very disappointing. Fine-grained parallelism of the PRAM algorithms require a high and frequently unreasonable amount of processors, in addition to the need of tightly synchronized lock-step operations among the many processors. It is therefore imperative to design models and algorithms in such a way that the theoretical complexity analysis matches the timings observed in actual implementations.

For geometric problems like visibility, convex hull, voronoi diagram, union of rectangles etc., which exhibit massive communication requirements, by using a coarse-grained model, Dehne [8, 9] has recently achieved considerable progress towards this goal. For many architectures, such results are a considerable improvement over existing methods. This applies in particular to previous fine-grained algorithms, even if they are (fine-grained) optimal. For example, it is impossible for fine-grained mesh algorithms, even optimal ones, to yield optimal speedups for ratios $n/p > O(1)$ (n = number of processors, p = number of data items) by applying the usual simulation method (also called "virtual processors" in many multicomputer operating systems). For $n/p > O(1)$, our methods are considerably faster. These algorithms are simple and easy to implement. The constants in the time complexity analysis are small. Except for a small (sometimes fixed) number of communication rounds, all other computation requires no communication. For coarse grained machines, this new method implies communication through few large messages rather than having many small messages. This is important for machines like the Intel iPSC where each message creates a considerable overhead. Dehne and his collaborators have implemented and tested prototypes of some of algorithms on a CM5 and iPSC/860, and obtained very fast running times which match nicely the theoretical analysis.

Dehne's coarse-grained model

We now present the main ideas of Dehne's scalable parallel computational model for coarse-grained distributed memory multicomputers.

Consider a problem of input size n and a parallel computer with p processors. Each processor gets thus n/p data items. PRAM algorithms often require $p = O(n)$. Thus the amount of data items to be placed in each processor is $n/p = O(1)$. Instead of this fined-grained model, we consider a more realistic case in which $p << n$, which is usually the case, even with current MPP (massively parallel processors). We want to design algorithms that are scalable, in the sense that they must be efficient for a wide

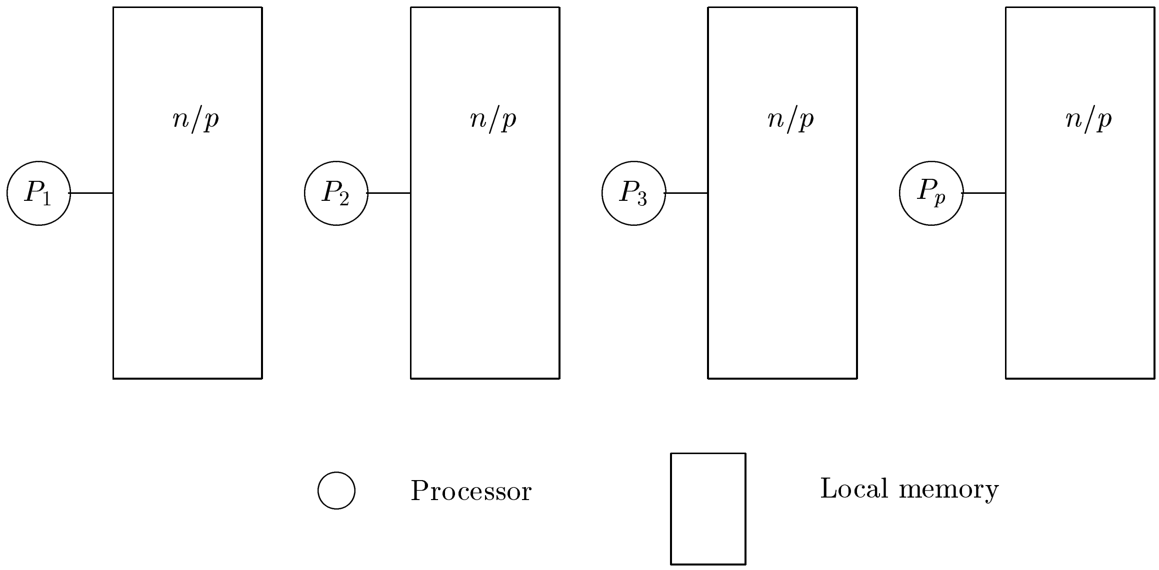

Figure 1: A distributed memory multiprocessor

range of n/p.

Consider Figure 1, where we have p processors $P_1, P_2, \ldots, P_p$, each with its local memory of size $O(n/p)$. The term coarse-grained is used here as to mean that the size of the local memory, $O(n/p)$, is "much larger" than $O(1)$. For instance, we may require $n/p \geq p$. These p processors are connected through some kind of interconnection network, e.g. mesh, hypercube, etc.

The basic idea is as follows. We first distribute the n input data items among the p processors each with n/p. This is done by a partitioning scheme. The algorithm to be designed consists of a repetition of two phases or rounds: a computation phase followed by a communication round. In the computation phase, we attempt to use the best known sequential algorithms in each processor that processes their data independently. In the communication round, we shift data around, with each processor sending out a total n/p data items and receiving a total of n/p data items. We express this by the following generic scheme.

repeat the following k times
    begin
        computation phase
        communication round
    end

The goal is to design algorithms which require a small number of communication rounds, i.e. with the smallest possible k. In each computation phase we use the best possible sequential algorithm. In many geometric algorithms presented in [8, 9], k is constant or $O(\log p)$. Since p is usually small as compared to n, the resulting algorithm will be highly efficient in practice. The studied geometric problems that can be solved as mentioned with k = constant or $O(\log p)$ include the following. Notice that efficient solution to some of these problems can be very important in Computer Graphics applications. The interested reader should refer to [8, 9] for details.

1. area of the union of rectangles
2. $3D$-maxima
3. $2D$-nearest neighbors of a point set
4. lower envelope of line segments in the plane (also known as the visibility problem)
5. $2D$-weighted dominance counting
6. multisearch on balanced search trees, segment tree construction, and multiple segment tree search.

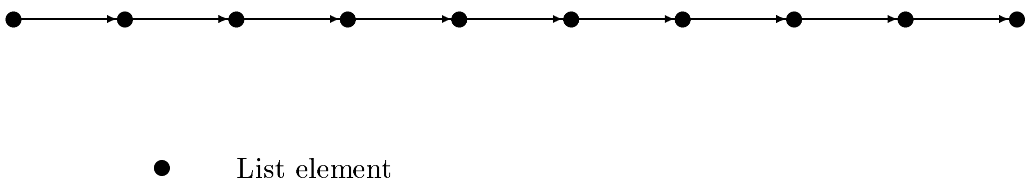

Figure 2: A linear linked list

In [10], we investigate the problem of list ranking. Consider a linked list of n elements (Figure 2). The problem of list ranking consists of finding the location or distance of each element in a linked list with respect to the end of the list. It appears as subproblem in several graph and tree problems. Therefore an efficient solution for list ranking has direct consequences in other applications. A trivial sequential algorithm solve this problem in $O(n)$ time by traversing the list. Several PRAM list ranking algorithms have been proposed, with $O(\log n)$ time complexity.

In our coarse-grained distributed memory model, we distribute the input n in the p processors, as shown in Figure 3.

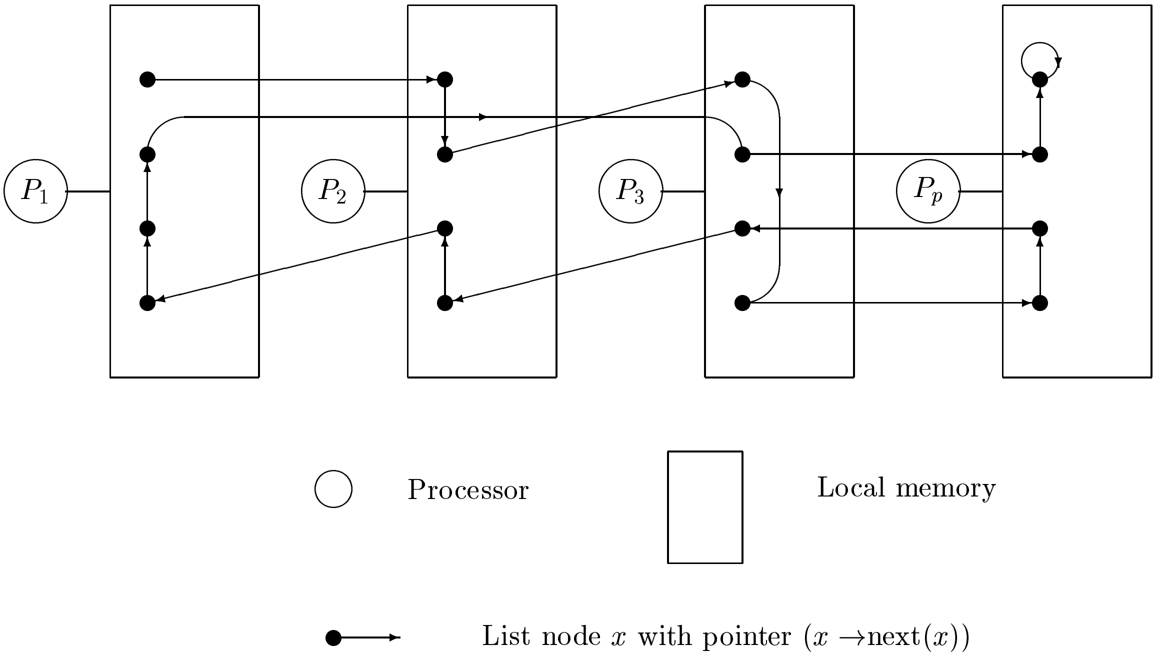

Figure 3: A linear linked list stored in a distributed memory multiprocessor

In [10] we present a scalable parallel algorithm for the n-element list ranking problem using p processors. Two versions of the algorithm are given. The first version requires, with high probability, $\log(3p) + \log\ln(n)$ communication rounds. The second version requires only $O(r \log p)$ communication rounds, with high probability, where $r < \ln^*(n)$ is an extremely small number.

It is not our purpose to give the details here (which can be obtained from [10]). We wish to emphasize, however, the importance of this coarse-grained computing model for distributed memory systems. If we manage to design appropriate algorithms with a small number of communication rounds, then it is very likely that the resulting algorithm will be efficient in practice.

6 Virtual shared memory systems

Shared memory systems have been popular due to many advantages and ease of their usage. Cache-coherent non-uniform memory access and non-coherent non-uniform memory access architectures have been proposed [1, 16]. The hardware version of shared memory incurs nevertheless high cost of implementation. Implementation of shared memory in software, also known as virtual shared memory, constitutes a viable alternative [4, 5].

Research in virtual shared memory systems are being carried out at the Universidade Federal de Rio de Janeiro on the NCP II machine (successor of the NCP I [2], and at the Universidade de São Paulo (IME/USP) [6, 15].

The NCP II project is financed by the FINEP/MCT (Financiadora de Estudos e Projetos – Research and Project Funding Agency) of the Ministry of Science and Technology and started at the beginning of 1995. This constitutes one of the few projects supported by FINEP in the area of design and development of high-performance computers in Brazil. Another project financed by FINEP is being carried out at LSI/USP (Laboratório de Sistemas Integráveis, Universidade de São Paulo). The FINEP support amounts to roughly four million dollars.

The research of virtual shared memory at the Universidade de São Paulo (IME/USP) is supported by a joint GMD/CNPq international cooperation program between Brazil and Germany. The counterpart in Germany is GMD FIRST, Berlin. One goal of the joint program, among other goals, is to develop performance prediction models for parallel programs, based on a virtual shared memory system (VOTE [6]) and runs on the GMD Manna machine. The following is an abstract extracted from [15].

The performance capabilities of virtually shared memory systems (VSM) still fails to seriously refute the impression that VSM seems to be a luxury, which provides conceptual simplicity at the expense of performance. The paper [15] is based on the VOTE VSM system running on top of a crossbar-based parallel computer system. The design of VOTE is a two layer approach, providing conceptual simplicity by means of a sequential consistency at default and, moreover, providing a framework of functions to allow performance enhancements. This enables the programmer to use a customized programming style addressing any combination of sequential consistency's conceptual simplicity and the maximum performance achieved through message passing.

To validate the approach, performance overheads are discussed. The paper [15] shows the possibility of developing analytical performance models. Such models are able to predict the execution time for a generic number of processors and for different architectures. They can be used for bottleneck identification and, in some cases, serve as a guide for software improvement for both the application and the VSM system.

Appendix A – Supercomputer systems

NEC SX3/12R

NEC SX3/12R (1 CPU, 2 pipes, 3.2 GFlops)

> Local: Instituto Nacional de Pesquisas Espaciais / Centro Meteorológico de Previsão e Tempo (INPE/CPTEC Weather Forecast Center)

Observation: This equipment occupied the 186th position on the TOP500 Supercomputer Sites, Tennessee/Mannheim report of 9 Nov. 1994. In the November, 11, 1995 report, it occupies the 420th position on the TOP500 list.

Cray Research EL98 and J-90

Cray Research EL98 with 6 processors
Cray Research J-90 with 8 processors (to be upgraded to 16 proc.)

Local: Universidade de São Paulo (USP).

Cray Research Y-MP2E

Cray Research Y-MP2E – 2 CPUs.

Local: Centro Nacional de Supercomputação, Universidade Federal do Rio Grande do Sul (CESUP/UFRGS)

Observation: one of the Brazilian National Supercomputer Centers (Centro Nacional de Processamento de Alto Desempenho – CENAPAD).

Cray Research EL94 and J-98

Cray Research EL94 and Cray J-98 (mid 95)

Local: Universidade Federal do Rio de Janeiro (COPPE/UFRJ)

IBM 9021-711

IBM 9021-711 with 1 VF.

Local: Universidade Estadual de Campinas

Observation: one of the Brazilian National Supercomputer Centers (Centro Nacional de Processamento de Alto Desempenho – CENAPAD).

IBM 3090/600

IBM 3090/600-6VF (6 proc.)

Local: Petrobrás, Rio de Janeiro.

IBM 9021-820

IBM 9021-820 (4 processors).

Local: Petrobrás, Rio de Janeiro.

IBM 3090

IBM 3090 with VF.

Local: Universidade Federal de Santa Maria (UFSM)

IBM SP-1

IBM SP-1 (16 procs.)

Local: Laboratório Nacional de Computação Científica, Conselho Nacional de Desenvolvimento Científico e Tecnológico (LNCC/CNPq).

Observation: one of the Brazilian National Supercomputer Centers (Centro Nacional de Processamento de Alto Desempenho – CENAPAD).

IBM SP-1

IBM SP-1 (8 procs.)

Local: Universidade Estadual de Campinas (UNICAMP).

Observation: one of the Brazilian National Supercomputer Centers (Centro Nacional de Processamento de Alto Desempenho – CENAPAD).

IBM SP-2

IBM SP-2 (4 procs.)

Local: Universidade Federal do Ceará (UFCE)

Observation: one of the Brazilian National Supercomputer Centers (Centro Nacional de Processamento de Alto Desempenho – CENAPAD).

IBM SP-2

IBM SP-2 (4 procs.)

Local: FUNCEME/Ceará.

Observation: one of the Brazilian National Supercomputer Centers (Centro Nacional de Processamento de Alto Desempenho – CENAPAD).

Intel iPSC-860 / 8

Intel hypercube iPSC-860 / 8

Local: Universidade Federal do Rio de Janeiro (COPPE/UFRJ)

Intel IPSC-2 / 16

Intel hypercube IPSC-2 / 16

> Local: Universidade de Brasília (Física/UnB).

Parsytec PowerXplorer

Parsytec PowerXplorer 16/32 (16 nodes each with a PowerPC 601 and T805 – 1.28 GFLOPS peak).

> Local: Universidade de São Paulo, Instituto de Matemática e Estatística (DCC/IME/USP).

Parsytec Multicluster/MP

Parsytec Multicluster/MP (Transputer).

> Local: Universidade Federal de Pernambuco (DI/UFPE).

Silicon Graphics Challenge

Silicon Graphics Challenge / 8 CPUs R4000

> Local: Universidade de São Paulo (LSI/USP).

Silicon Graphics Power Challenge

Silicon Graphics Power Challenge (6 procs.)

> Local: Instituto de Estudos do Mar/Rio de Janeiro

Observation: This equipment occupied the 392th position on the TOP500 Supercomputer Sites of the TOP500 Tennessee/Mannheim report on Nov. 9, 1994.

Silicon Graphics / 2 CPUs R800

Silicon Graphics / 2 CPUs R800

> Local: Centro Tecnológico Aeroespacial - Instituto de Estudos Avançados (CTA/IEAv)

Silicon Graphics Onyx

Silicon Graphics Onyx (4 procs. - shared memory)

> Local: Centro Nacional de Supercomputação, Universidade Federal do Rio Grande do Sul (CESUP/UFRGS)

NCP-1

Hypercube NCP-1 (8 nodes each with a T800 and Intel i860).

Local: Universidade Federal do Rio de Janeiro (COPPE/UFRJ)

Wavetracer

Wavetracer / SIMD

Local: Universidade Federal de Minas Gerais (DCC/UFMG)

Appendix B – World Wide Web sites

We list below some of the WWW sites related to Latin American universities and R&D related institutions.

Table 8

Argentina (top level domain)	http://www.ar:70/
Universidad de Chile	http://www.dcc.uchile.cl/
Costa Ricas Research Network	http://ns.cr/
Ecuador: Universidade San Francisco de Quito	http://mail.usfq.edu.ec/
Mexico (Info)	http://info.pue.udlap.mx/
Pont. Univ. Catolica Rio de Janeiro (PUC/RJ)	http://www.puc-rio.br/
Univ. Federal do Rio de Janeiro (COPPE/UFRJ)	http://guarani.cos.ufrj.br:8000/
Univ. Federal de Minas Gerais	http://dcc.ufmg.br/
Univ. Federal do Rio Grande do Sul (Inst. Inf.)	http://tucano.inf.ufrgs.br/
Univ. Federal do Rio Grande do Sul (CESUP)	http://www.cesup.ufrgs.br/
Univ. Federal de Santa Catarina	http://www.inf.ufsc.br/
Universidade Estadual de Campinas	http://www.unicamp.br/
Universidade de Sao Paulo	http://www.usp.br/
	http://www.lsi.usp.br/
FAPESP Fund. de Amparo a Pesq. Est. S. Paulo	http://www.fapesp.br/
CNPq Conselho Nacional de Des. Cient. e Tecnol.	http://www.cnpq.br/
CITE/CNPQ (Softex 2000, ProTem, RNP)	http://www-cite.cnpq.br/
INPE - Inst. Nac. de Pesquisas Espaciais	http://www.inpe.br/
LNCC - Lab. Nac. Computacao Cient.	http://www.lncc.br/

7 Conclusion

By observing the evolution of the TOP500 supercomputers in terms of architecture during the past few years, we notice an increasingly higher number of MPP (massively parallel processors) systems. We can be quite certain that Parallel Computing will play a crucial role in high-performance computing in the next decade.

Although there are very few supercomputers on the TOP500 list located in Latin America, we notice a considerable number of "smaller" supercomputer sites in Latin America, as shown in Appendix A.

The role of the federal government in establishing public scientific policy is very important. A notable example is the Brazilian Policy on Science and Technology of CNPq/MCT (Conselho Nacional de Desenvolvimento Científico and Tecnológico – National Science and Technology Development Agency) and FINEP/MCT (Financiadora de Estudos e Projetos – Research and Project Funding Agency), both of the Ministry of Science and Technology. These agencies determines the scientific and technological policies and provide the necessary funding. In addition to these two, there is still CAPES/MEC (Fundação Coordenação de Aperfeiçoamento de Pessoal de Nível Superior) of the Ministry of Education that contributes to the formation of human resources by providing scholarships to graduate students to pursue their degrees, both in Brazil and abroad.

We examined the important role of funding resulting from international cooperation programs. To obtain such fundings some of the usual requirements include the presence of international partners, including industrial partners. The projects are usually application-oriented and should produce measurable deliverables that benefit the society. Though high-performance computing is costly, it can be crucial to solve some the very pressing problems that can benefit mankind (environment analysis and planning, meteorology, etc.).

In more technical terms, we have examined the importance of defining new models of computation and the design of scalable parallel algorithms. Such algorithms not only should be efficient in theory, but principally in practical implementations. We also called our attention to the research of virtual shared memory that unites conceptual simplicity and the maximum performance.

8 Acknowledgment

This research was supported by FAPESP (Fundação de Amparo à Pesquisa do Estado de São Paulo) - Proc. No. 95/0767-0 and 95/6159-1, CNPq (Conselho Nacional de Desenvolvimento Científico e Tecnológico) - Proc. No. 523112/94-7 and PROTEM-2-TCPAC Proc. No. 680060/94-4, Commission of the European Communities through Project ITDC-207, and UNIDO (United Nations Industrial Development Organization). I would like to thank Prof. Boleslaw K. Szymanski for inviting me to prepare this paper and for his comments. I am also grateful to Prof. João Paulo Kitajima from the Universidade de Brasília who compiled some of the material used in this paper.

References

[1] Agarwal, A. et al. "The mit alewife machine: architecture and performance". *Proceedings 22nd International Symposium on Computer Architecture.* June 1995.

[2] Amorim, C. L. et al. "O sistema de computação paralela NCP I". *Anais do 7.o Simpósio Brasileiro de Arquitetura de Computadores e Processamento Paralelo.* July 1993, pp. 89 –99.

[3] Baeza-Yates, R. A., Fuller, D. A., Pino, J. A. and Goodman, S. E. "Computing in Chile: the Jaguar of the Pacific Rim?". *Communications of the ACM*, Vol. 38, No. 9, pp. 23 – 28, 1995.

[4] Bershad, B. N., Zekauskas, M. J. and Sawdon, W. A. "The midway distributed shared memory system". *Proceedings of COMPCON Spring 93.* 1993.

[5] Cordsen, J. "Basing virtually shared memory on a family of consistency models". *Proceedings of the International Workshop on Support for Large Scale Shared Memory Architectures*, pp. 58 – 72, April 1994.

[6] Cordsen, J. and Gubitoso, M. D. "Performance considerations in a virtualyl shared memory system". Technical Report – PLENA project. GMD - FIRST and Universidade de São Paulo. September, 1995.

[7] Dedrick, J. L. "Little Engines that Could: Computing in Small Energetic Countries". *Communications of the ACM*, Vol. 38, No. 5, pp. 21 – 26, 1995.

[8] Dehne, F., Fabri, A. and Rau-Chaplin, A. "Scalable parallel computational geometry for coarse grained multicomputers," in *Proc. ACM Symposium on Computational Geometry*, pp. 298-307, 1993.

[9] Dehne, F., Kenyon, C. and Fabri, A. "Scalable And Architecture Independent Parallel Geometric Algorithms With High Probability Optimal Time," in *Proc. IEEE Symposium on Parallel and Distributed Processing*, Dallas, Tx, pp. 586-593, 1994.

[10] Dehne, F. and Song, S. W. "Scalable Parallel Pointer Jumping For Distributed Memory Multiprocessors, And Applications". Technical Report, Universidade de São Paulo, Instituto de Matemática e Estatística, Departamento de Ciência da Computação. In preparation.

[11] Dongarra, J. J., Meuer, H. W. and Strohmaier, E. "TOP500 Supercomputer Sites". Technical Report RUM 33/93. Computing Center, University of Mannheim. November 11, 1993.

[12] Dongarra, J. J., Meuer, H. W. and Strohmaier, E. "TOP500 Supercomputer Sites". Technical Report RUM 42/95. Computing Center, University of Mannheim. November 10, 1995.

[13] EITO European Information Technology Observatory, 1994.

[14] Estivill-Castro, V. "Computer Science in Mexico". *IEEE Computer*, Vol. 28, No. 8, pp. 56 – 62, 1995.

[15] Gubitoso, M. D. and Cordsen, J. "Performance considerations in VOTE for PEACE". Technical Report RT-MAC-9514, Universidade de São Paulo, Instituto de Matemática e Estatística, Departamento de Ciência da Computação. November 1995.

[16] Kuskin, J. et al. "The flash multiprocessor". *Proceedings 21st International Symposium on Computer Architectures.* April 1994.

[17] La Rovere, R. L. and Goodman, S. E. "Computing in the Brazilian Amazon". *Communications of the ACM*, Vol. 35, No. 4, pp. 21 – 24, 1992.

[18] Lux, T., Sydow, A. and Schmidt, M. "An Air Pollution Model Used for the Region of Berlin". *Proc. of 7th International Conference on Computer Methods for Partial Differential Equations*, June 1992.

[19] Rosenthal, D. and Meira, S. L. M. *Os Primeiros 15 Anos da Poltica Nacional de Informática: o Paradigma e sua Implementação.* CNPq/ProTem-CC – Conselho Nacional de Desenvolvimento Científico e Tecnológico/Programa Temático em Ciência da Computação. 285 p. 1995.

[20] Softex 2000, CITE/CNPq (Coordination of Technology Innovation, Conselho Nacional de Desenvolvimento Científico e Tecnolgico) – http://www-cite.cnpq.br/