Parallel Computing: New Opportunities for Transportation Research

T.G. Crainic[1], M. Gendreau[1], G. Laporte[1] and C. Lardinois[2]

This paper describes a number of applications of parallel computing in the field of transportation, as well as some research avenues derived from this technique. The impact of parallel computing on the study of several classical operations research problems arising in transportation is examined, and a number of industrial and commercial applications where parallel computing could be relevant are described.

INTRODUCTION

In the last ten to fifteen years, factors such as urban sprawl, deregulation and the dismantling of international trade barriers, the fall in the price of raw materials and the rise of competition from abroad have caused profound changes in the field of transportation. As a consequence, the working environment prevailing at the turn of the 21st century is quite different from that of the 1970s and 1980s. At the same time, tremendous progress in electronics, computing and telecommunications has opened up new horizons for today's technicians and decision-makers by creating a steady stream of new technological alternatives. Major changes have been felt in different areas of transportation application. Real time control of urban traffic using computers, in-vehicle communication and guidance systems, satellite-based systems to locate and identify trucks and trains operating in a given area, and simultaneous day-to-day or even hour-to-hour operations of carriers and multimodal logistical chains through an exchange of data and computerized documents are a few examples of these modifications.

These technological developments frequently generate large and intrinsically complex problems requiring very fast process of data; in addition, they are often accompanied by considerable uncertainty and randomness. These requirements, in turn, make increasingly greater demands on technology, especially with respect to computer processing power. Parallel computing appears to be a partial answer to these problems, as well as a source of new scientific challenges.

The purpose of this paper is to review possible applications of parallel computing in transportation as well as some research perspectives opened up by this technology. In the following section, parallel computing is presented as a natural complement to conventional operations research methods used to develop transportation systems planning and management tools. Then, the relationship

1. Centre de Recherche sur les Transports, Université de Montréal, C.P. 6128, Succursale Centre-Ville, Montréal, Canada, H3C 3J7.
2. Ecole de Technologie Supérieure, 4750 Avenue Henri-Julien, Montréal, Canada, H2T 2C8.

Scientia Iranica, Vol. 2, No. 3, © Sharif University of Technology, October 1995

between parallel computing and the theoretical and methodological foundations of transportation research is discussed, and the contribution of parallel computing to the study of several frequently encountered problems is examined. Three application areas in which parallel computing can be justified from an industrial and commercial point of view are then presented.

THE NEED FOR COMPUTING POWER AND PARALLEL COMPUTING

Planning and management tasks have always been varied and complex in the transportation industry. However, for many years, and especially in the last ten to fifteen years, the transportation field has had access to a number of software tools and other analytic and decision-making support systems developed using results from computer science and operations research. The EMME/2 [1], STAN [2] and HASTUS [3] systems, designed as planning and management tools, are just a few examples of the scientific, technical and commercial success brought about by jointly applying computing and operations research to transportation problems.

While this winning combination continues to provide the most dependable aid towards efficiency, productivity and competitiveness in transportation, it is clear that available means and tools are not always up to the task. Given the current state of science and technology, operations research methods as well as microcomputers and work station architecture geared to serial processing often lack the computational power required in many advanced transportation fields. If more conventional requirements of flexibility, portability, user-friendliness and, above all, cost-effectiveness are added in, then the appropriateness of both small computers and supercomputers becomes questionable as well.

The operations research problems generated by the new systems just mentioned and described in more detail later in this paper have a common and essential characteristic that lends itself particularly well to the use of parallel-architecture computers. All of these problems can be modeled as network problems and are decomposable, frequently in many different ways (for example, by groups of vehicles, origin-destination pairs and so on) into subproblems of smaller size or complexity. These two key ideas, networks and decomposability, combined with the need for computational power suggest strongly that the potential of parallelism should be explored, since they raise the possibility of simultaneously processing different parts or components of the initial network problem on a number of processors operating in parallel.

The rise of parallel computing, an important and inescapable trend in the development of computer architectures, is in fact a direct consequence of the physical characteristics of the microelectronic substructure. Current "serial" architectures are now approaching performance limits that probably can only be circumvented through the use of many processing sites in simultaneous and cooperative operation, that is, through the use of parallel computing. The need to determine conditions for effective use of this new technology has given rise to a research field generally viewed as highly strategic, thus, posing a number of new challenges. In particular, many issues are raised by the various forms that parallel computing can take, ranging from a network of powerful systems linked together relatively loosely (such as the MIMD, Multiple Instruction Stream or Multiple Data Stream approach) to those made up of hundreds of thousands of smaller processors equipped with dense and very fast means of communication (such as the SIMD or Single Instruction Stream Multiple Data Stream approach), also called massive parallelism. Experimental machines combining both approaches are starting to make their appearance.

The high cost of this equipment raises the issue of cost-effectiveness. A good deal of research and development work in parallel computing is currently under way which generally uses expensive equipment, and therefore, can only be justified in very well funded areas such as defense and aerospace research. Fortunately, new technology based on networks of workstations and multiprocessor servers has lowered the costs substantially while maintaining a very high level of efficiency and power. In addition to being an excellent solution to the cost-efficiency tradeoff, this technology also has the considerable advantage of supporting modularity by allowing the construction of the systems capable of growing gradually and increasing in power as requirements dictate.

All areas in which computers are applied will eventually benefit from parallelism and in fact this new technology is already becoming commercially available. While widespread use of this technology is still a long way off, it is easy to predict that technological progress and increasing demand will bring certain forms of moderately priced parallel computing to the marketplace and that parallelism will become more prevalent. This scenario, seen as very realistic by most if not all experts, constitutes an exceptional opportunity for research and development in the transportation industry since parallel computing is particularly well suited to the requirements and technical characteristics typical of a good deal of advanced transportation research.

THEORETICAL AND METHODOLOGICAL FOUNDATIONS

Concurrently with developments in computer technology, the vast field of quantitative methods, particularly operations research applied to transportation systems planning and management, has grown steadily in the last forty years. One aspect of this joint evolution has been the identification of a set of increasingly complex base problems, which have then been tackled using a battery of increasingly efficient mathematical models and algorithms. These problems, models and algorithms make up the theoretical and methodological foundations of operations research in transportation and have provided the framework for development of computerized tools for transportation planning. The base problems of interest here may be discussed along three dimensions:

Mathematical Modelling

Parallel approaches can be expected to allow a more detailed description of the problems at hand. In some instances, this could lead to radically different models where problems would be stated in a more natural way instead of having to rely on somewhat unsatisfactory aggregations, which are currently used because of computational limitations.

Algorithm Processing

Since parallelism is based on the idea of distributing algorithmic tasks among a number of processors, these solution algorithms must be designed and structured to reflect this task distribution and the type of parallel architecture planned for use.

Computer Systems Management and Implementation

Model and algorithm design and efficiency in a parallel environment are strongly affected by each of the following factors: the physical organization of the processor network (spatial organization, physical or logical connections, etc.); the network's functional organization, or the roles defined for each of the processors (for example, hierarchical—also called master/slave—or collegical organization) and the type of communication (synchronous or asynchronous) between processors; and algorithm coding itself.

To illustrate this, three broad classes of problems and general methods that are particularly well suited to parallel computing are discussed: 1) static network flow problems; 2) problems that explicitly incorporate the dynamic and probabilistic aspects often typical of transport systems and operations; and 3) more general solution methods applicable to a variety of frequently occurring transportation problems.

STATIC NETWORK FLOW PROBLEMS

Combined under this heading are five well defined base problems and fundamental operations research models in transportation.

Shortest Path Problems

The goal of this class of problems is to find the shortest path or paths—in terms of distance, time or any other generalized cost criterion—between various nodes of a transportation network; from one node to another, from one node to all the others, or between all pairs of nodes. These fundamental problems appear in a number of different forms in all areas of both passenger and freight transportation. Here, of main interests are the very large-scale problems where parallel computation techniques can be applied to the analysis of various algorithms, particularly algorithms based on graph decomposition [4,5], guided search [6] and, more recently, dual "auction" methods [7,8].

Multicommodity, Minimum Cost Flow Problem with Linear Arc Costs

For a given transportation network with arc capacities, unit transportation costs and specified O-D matrices describing, for example, demand for transportation of various products, the objective of problems in this class is to simultaneously assign O-D flows in order to minimize total transportation cost while respecting capacity constraints. Problems of this type, especially large ones, are difficult to solve because of "binding" constraints on the arcs of the network. Whenever possible, these constraints are relaxed to yield, in most cases, a multicommodity flow problem with convex costs [9,10]. Otherwise, decomposition methods must be used. An initial parallel approach would be to allocate one or more of the subproblems resulting from this decomposition to each processor, while a "master" processor distributes and coordinates the tasks [11]. This promising approach does not seem to have been studied in depth while the literature has devoted a good deal of attention to very fine-grained parallelization (for example [12]) on very powerful machines (massively parallel and extremely costly).

Multicommodity Network Flow Problems with Convex Costs (Minimization of Total Cost)

Here, congestion is formalized by representing arc costs using convex functions. Such problems arise in several transportation models such as freight transportation, where flows of several types of commodities use the same network. These are also encountered in other areas, such as traffic assignment in telecommunications networks. The natural decomposition by commodity has led to development of Gauss-Seidel solution methods that decompose by commodity block [13]. A "two-level masterslave" approach to parallel computing is possible, where each subproblem becomes a single-commodity flow problem with convex costs, which is then solved by the linear approximation method [13]. If there are enough processors, this method can be effectively parallelized by arranging processors hierarchically [14], or by adequately transforming the model [15].

Static Assignment Problems (Equilibrium Models)

This includes flow problems mainly applicable to urban passenger transportation (using public transit or private cars). In this setting, one does not seek "system-optimal" flows, but rather a "user-optimum" in which each user attempts to minimize his or her own transportation costs. The static case, considered here, models the average state of the system over a given period, for example in the morning rush hour. Formulations of these problems in the space of arc flows have been studied widely [16], and conventional implementations use the linear approximation method [1]. A parallel approach to this method, in which arc processing is distributed among different processors, has already been implemented successfully on a transputer network with no impact on convergence rate [14]. Other parallel approaches have also been shown to be effective for problems with about ten O-D pairs [17]. Parallelization of algorithms that converge more rapidly than the linear approximation method, such as decomposition (Jacobi and Gauss-Seidel) and second-order (Newton and quasi-Newton) methods, could be considered for solving the problem formulated in the space of path flows [18]. Also conceivable is a parallel approach to the decomposition method that distributes blocks (O-D pairs) among the various processors, as well as an implementation of an asynchronous version of this method [19]. An approach to parallelization by network partitioning is possible for other methods (gradient, Newton, etc.), although these approaches do not seem to have received any attention in the literature.

Transportation Network Design Problems

This very broad class of problems covers all situations in which a transportation network (physical infrastructure or service network) must be defined to satisfy transportation demand usually expressed in the form of one or many O-D matrices. The objective function to be optimized generally includes terms for cost of construction in the case of physical infrastructure, operation and use of the defined network in the case of services and, where appropriate, other terms to describe network performance. For example, one could consider the measure of pollution resulting from use of an urban highway network, or measures of quality of freight service. The problem is to find an acceptable trade-off between these often contradictory cost components.

Two large subclasses of such problems can be distinguished: first are those in which the organization or business firm carrying out the design has control over the itineraries of the transported "objects" as is ordinarily the case in freight transportation; the second are those in which users travel on the network in such a way as to minimize their own transportation cost, the classic example of this being a highway network in which drivers choose itineraries independently of one another. In both cases, direct mathematical solution of the problem is difficult from an algorithmic point of view for a number of reasons: networks to be solved are almost always very large; simply evaluating a possible network design from the point of view of the objective to be optimized involves the solution of a network flow subproblem to determine itineraries on the network; more often than not problems are highly combinatorial in nature. This highly combinatorial aspect of these problems implies use of branch-and-bound [20-22] or of powerful heuristics.

Because of these problems' high degree of complexity, various parallel approaches are conceivable. One approach would be to parallelize the search strategy for finding a good exact or approximate solution. In this approach, many processors would consider many possible designs simultaneously, with sequential computations required for evaluating a single design falling to a single processor. Alternatively, one could use a sequential search strategy in which evaluation of each design is parallelized using parallel subproblem algorithms. Finally, if enough processors are available, the two approaches could be combined by parallelizing both search strategy and subproblem solution.

DYNAMIC NETWORK FLOW PROBLEMS

In dynamic problems and models, evolution of the transportation system over time is taken explicitly into account. In many cases, the inherent uncertainty of particular aspects of these systems can be considered as well. It is clear that formulating and solving such problems requires an adequate understanding of simpler static and deterministic problems. In fact, the high level of detail and complexity brought into these models by the addition of dynamic and stochastic dimensions indicates that parallelism will be the preferred approach and, at the same time, will enlarge the field of applications.

Problems of Planning and Management on Dynamic Stochastic Networks

This class of problems is extremely varied and far-reaching, since it pertains to all situations in which persons or objects move in a network for which particular parameters, such as arc travel time, capacity or demand, are not known with certainty but are described using random variables with known probability distributions. A typical example of such a problem is that of operational management of a fleet of containers over a multi-period planning horizon, in which empty containers must be allocated and repositioned to satisfy known present demands and uncertain future demands, while minimizing container transport costs and penalties for unmet requirements (see [23] for further details). These problems are usually formulated as stochastic programs, in general fairly difficult to solve using conventional approaches, especially if the network to be solved is large. On the other hand, parallel computing methods open the way for efficient algorithms to solve these problems. Three particularly promising approaches are mentioned here.

The first is based on an iterative stochastic programming algorithm that solves sets of deterministic subproblems corresponding to all possible realizations of the random network parameters [24]. This algorithm can be greatly speeded up by parallelizing the phases of subproblem solution and by assigning solution of subproblems corresponding to one or more parameter scenarios to individual processors. Subproblems usually take the form of minimum cost flow problems with linear or convex costs. This approach is similar to that of Mulvey and Vladimirou [25,26] developed to analyze financial phenomena for which parallel solution has generally required massively parallel machines [12]. The second approach is based on a partition of the network into several parts. At each stage of the selected solution algorithm, the network is cut up into large, nearly independent pieces; calculations for each piece are performed on a separate processor. Finally, the dynamic structure of the problem could conceivably be exploited by decomposing it by period. In such a decomposition, each processor would be responsible for a specific period and would do all calculations relating to that period. Such an approach does not seem to have ever been implemented, but there is room for very interesting theoretical and algorithmic developments, comparable to direct applications of Benders' decomposition [27].

Dynamic Assignment Problems

These are traffic and transit assignment problems in which interest is focussed on the dynamic aspect of phenomena under study, rather than a single and typical state of the system in a given period. Several models have appeared in the scientific literature. Some of these assign flows on a space-time hypernetwork constructed either explicitly [28-30] or implicitly [31,32]. Other models, based on user learning hypotheses, are more like simulation than optimization, in that for each iteration they suppose that users daily reassess their paths based on their knowledge of the state of the network in previous days. Sawack and Thompson [33] discuss the case of public transit; Drissi-Kaïtouni [34] and Drissi-Kaïtouni and Gendreau [35] discuss the case of private transportation.

In the case of the first type of model, serial and parallel algorithms developed for the static case could be adapted to the space-time hypernetwork. In the case of the second type of model, the network could be subdivided and distributed among the various processors. At each time period, individual processors handle flow dynamics on the corresponding subnetwork, and satisfactory transfer of flows at subnetwork boundaries from a processor to its neighbours is handled by a communications process.

SOLUTION METHODS

Solution methods pertinent to several large classes of problems, particularly those with combinatorial or stochastic nature are discussed in this section. Three classes of methods are of particular interest.

Branch-and-Bound Methods

For many combinatorial optimization problems, the only practical exact solution approach is branch-and-bound, in which a tree structure is imposed progressively on the space of possible solutions. Since the solution tree can grow quite large, these methods are generally slow. A parallel approach offers potential for fast and exact solutions to problems for which, up to now, approximate solutions have had to suffice [36]. A number of studies have examined the use of parallelism in branch-and-bound methods [21,22,37]. For example, the travelling salesman problem, heavily combinatorial and very difficult to solve exactly, is a fundamental problem in operations research. Pruul [38] applied MIMD methods to solve this problem, while Kindervater and Trienekens [31] and Trienekens [39] showed that a master-slave approach can be effective when the length of time each processor must work is greater than the inter-processor communication time, and may even be better than collegial asynchronous methods [40]. Specific methods have also been developed to take advantage of available machine architectures [10]. Examples are the distributed environment discussed by Finkel and Manber [41] in which several limited-capacity processors interrelate and cooperate by exchanging messages, and the hypercube discussed by Pardalos and Rodgers [42]. Branch-and-bound methods have also been successfully parallelized in the context of network design [37]. In particular, it has been shown that a transputer network's physical and communications topology has a profound impact on algorithmic performance [20].

Other applications of parallel branch-and-bound methods to stochastic and network design problems are conceivable. Thus, several other approaches to task distribution, as well as combination with parallel tabu search methods, could be developed to take advantage of the various structures and mechanisms for transputer network control and communication.

Tabu Search Methods

When difficult combinatorial problems have to be solved, it is often not possible or desirable to determine an exact solution. In such situations, various heuristic approaches can be used to find a good approximate solution. Tabu search methods are one of the most recent of these approaches. They explore the space of possible solutions iteratively; each iteration retains information on solutions visited in preceding iterations, to prevent the search from cycling [43,44]. Recent studies [45,46] have shown the enormous potential of these methods for both pure and mixed integer programming problems. These methods can be parallelized advantageously by searching the solution space and starting different search threads at different points. When and how processors communicate among themselves to exchange information, what information is exchanged, and how it is processed to increase the global knowledge and to reorient the search toward more promising parts of the solution space, make a wide spectrum of parallel search strategies [47–49].

Artificial Intelligence Methods

These methods all borrow concepts from theoretical computer science, particularly artificial intelligence approaches such as heuristic methods and neural networks. In heuristic routing methods, applied most frequently in freight transportation but also used in passenger transportation, vehicle routes are constructed iteratively by inserting locations that have not yet been visited [50]. Two such approaches are of interest here. First, a number of conventional heuristics developed some years ago could be adapted to the new parallel architectures [50,51]. In addition, some new methodologies stand to make particularly good use of parallelism. For example, a number of artificial intelligence techniques, such as bundle search, include parallel exploration of distinct regions of the search space [52,53]. In fact, these methods construct several distinct solutions in parallel by simultaneously considering several insertion possibilities. In this way, they naturally suggest an implementation using distributed architecture. One of the uses of neural networks is to allow machines to learn to reproduce the decision-making process of an expert, like a vehicle dispatcher, using examples of previous decisions [54-56]. The elementary components of such networks are neurons, which essentially operate locally, receiving signals from adjacent neurons, processing these signals locally and producing a response that is then communicated back to adjacent neurons. In such a distributed mode of operation, a parallel implementation could be constructed in which each neuron is supported by a separate processor. Clearly, such an implementation would lead to substantially faster learning, which is now very slow using sequential architectures.

ADVANCED TRANSPORTATION APPLICATIONS

The way in which unprecedented socio-economic requirements have led to the introduction of parallel computing into the transportation field is already discussed. Combined with the decomposability of the relevant operations research methods and models, these very real requirements seem to demand the use of parallel architectures. The present section, devoted to applications, will reveal the genuine industrial and commercial justification for parallel computing by discussing three advanced application areas. As in the previous discussion, each of these areas combines important research challenges with real industrial and commercial potential.

DYNAMIC VEHICLE FLEET MANAGEMENT IN REAL OR QUASI-REAL TIME

Dynamic fleet management problems are encountered frequently and in a variety of fairly different settings.

Dynamic Urban Vehicle Dispatching

This problem is relevant, for example, to taxi firms, carrier services, companies or agencies providing handicapped transportation, emergency services and repair services, as well as local pick-up and delivery services for intercity freight transportation. The problem is to allocate, in real or quasi-real time, vehicles to requests from clients spread over a given territory and requiring fast, if not immediate service. In general, such problems involve a good deal of randomness. For example, the location or time of certain requests are not known in advance or the visiting time at the request location may be highly variable. In reality, responsibility for dispatching falls to an individual who, through experience, manages to provide the requested services. Due to the complexity of this process, it has become truly essential to design and develop decision support systems. While a number of such systems already exist, for the most part these are only information systems to process requests, manage client information and update vehicle information such as location and current activity. They do not actually provide a tool that suggests alternative strategies for assigning vehicles to requests; such tools fall within the domain of operations research problems and methods applied to transportation.

Vehicle dispatching problems can be classified into two main categories. Typical of the first of these are emergency services, in which waiting vehicles must be situated throughout the catchment area to ensure complete and homogenous coverage. For each emergency call, a waiting vehicle must be selected and the most rapid itinerary to the call location and then to the final destination must be identified. Essentially, this first class of situations involves the shortest paths on large and dense networks.

Courier, pickup and delivery services are typical of the second category. The problem in this case is to essentially establish routes for vehicles, with the possibility of dynamic adjustment—that is, real-time route adjustments made possible by in-vehicle communications systems—to react to random fluctuations in the state of the system, such as the arrival of new requests.

The principal problems identified in the previous discussion to justify the use of parallel computing (large and dense networks, the need for quick processing and the necessity to include dynamic and stochastic dimensions) are present in all these situations. Since the decision processes of expert dispatchers in real-time dispatching are difficult to model mathematically, the method of neural networks seems particularly appropriate. Techniques by which a computerized system can learn to reproduce a decision-making process from examples of previous decisions appears very promising [54].

Real-Time Management and Control of Bus Services

During transit system operations, many factors such as accidents, weather conditions, vehicle delays or a sudden rush of passengers at certain stops, may perturb a planned bus schedule and cause service quality to deteriorate unacceptably. When this happens, schedules or operations must be rearranged as soon as possible to put operations back in order. For this reason, a number of transit agencies have considered adopting automatic vehicle location and control systems. These systems are useful for locating operating buses at any given moment with relative precision, and for communicating with buses to transmit orders to correct unacceptable situations.

For such systems to be effective, models and algorithms must exist that can identify in real-time the probable consequences of service perturbations and corrective actions, and thus choose the action most appropriate for achievment of the stated objectives. In view of the heavy element of randomness in transit system operations, it seems natural to formulate the real-time control problem for such a system as a problem of flow on a dynamic stochastic network. However, even the approximate real-time solution of reasonably sized problems of this type requires specialized algorithms to exploit, for example, the potential of parallelism.

Dynamic Management of Interurban Transportation Systems

Falling under this heading is a class of real or quasi-real time vehicle dispatching and fleet management problems having in common the fact that the "transportation" function is carried out on an interurban network. This class includes such problems as dynamic vehicle allocation, dynamic schedule adjustment and dynamic evaluation of transportation demand and service adjustment. The main challenges posed here are the large size of the problems considered, intertemporal relations between decisions made now and in future periods, uncertainty regarding demand service time and trade-offs between economic and service quality objectives.

Freight Marketplaces

This area is concerned with the establishment and management of electronic data interchanges and management systems to facilitate and coordinate the convergence of supply and demand in multimodal freight transportation. What must be determined, in particular, is the best multimodal route for each service request, based on request-specific criteria and in such a way that guarantees fair treatment of all firms and agencies doing business within the framework of the freight marketplace.

INTELLIGENT TRANSPORTATION SYSTEMS (ITS)

In the past few years, the United States, Europe and Japan have launched vast R& D programs to develop a group of software and hardware systems known as Intelligent Transportation Systems (ITS) also referred to as Intelligent Vehicle-Highway Systems (IVHS). These systems call on various industrial sectors and technologies, including information processing, computing, telecommunications, control and electronics. Their goal is to reduce automobile congestion and pollution, enhance mobility, improve the productivity of public and private vehicle fleets and increase road safety. IVHS America (USA), DRIVE and PROMETHEUS (Europe), RACS and AMTICS (Japan) are some of these well-known undertakings. Canada is also gradually embarking on ITS projects [1].

The following systems and services (some of which are already being developed or tested, at least in part) illustrate the ITS concept:

- Information and communication systems linked to control centers that advise drivers about current and expected traffic conditions, road hazards, etc., and that suggest alternate routes.

- Transportation management systems that flexibly adjust road usage, speed limits, traffic signals and roadway access based on current traffic conditions.
- On-board electronics and guidance systems to help drivers plan and follow safe and efficient routes.
- Tracking and dispatching systems for commercial fleets (trucks, delivery vans), transit operators and emergency vehicles to offer productive and flexible demand-responsive services.

The advent of such advanced systems is driven by current and anticipated progress in various electronic and communications subsystems and technologies for: real-time traffic counting, data communication, automatic positioning, in-vehicle driver interfaces, etc. However, if they are actually to bring about productivity, efficiency and safety gains in large, complex transportation systems, these ITS technologies must also integrate a great deal of "intelligence" in accordance with their name. Whatever the sophistication of the hardware, it remains useless if one does not know what to do with the huge amount of information which will be collected and made available from various sources. One needs advanced capabilities to model, operate and monitor—in real or quasi-real time—transportation systems involving complex network-based dynamic interactions, often accompanied by considerable uncertainty and randomness. This level of sophistication in the "methodological and intelligent part" of ITS will make the difference between successful systems and others.

Dynamics, real-time, randomness, problem size and complexity are the key issues to be addressed when capturing, processing and communicating information within ITS. Hence, the need is felt for very fast processing of information and quick responses, which in turn makes increasingly greater demands on technology, especially with regard to computer power. In this respect, it is clear that standard sequential computers are not up to the task, and that other types of architectures must be considered. Vector computers and massively parallel machines have been used in some ITS tests, but the high cost of such equipment makes it prohibitive in most real-world operational conditions. Smaller coarse grain parallel computers that are affordable by most organizations and may grow with users' needs, appear to be the most promising technology.

TOOLS FOR TRAINING, EXPERIMENTATION AND TECHNOLOGY TRANSFER

Three questions are relevant here: 1) How can individuals be trained for complex planning and management work in transportation firms or agencies? 2) How can new strategies or techniques be evaluated, without having to carry out in situ experiments? 3) How can the transportation industry become aware of this and be encouraged to use new methods and tools for analysis and decision support? These three questions raise the same challenge: technological modernization and the development of specialized skills and knowledge in the area of transportation planning and management.

To provide a concrete and scientific answer to these three questions, researchers have recently begun to design and develop simulation games for use as 1) training instruments for transportation planning and management; 2) experimental frameworks for analyzing and evaluating different scenarios, strategies or techniques; and 3) mechanism for technology transfer between the university and industry, that is, as a mean of encouraging transportation companies and agencies to become familiar with the newest analysis, planning and management methods. These simulation games are designed as business games that imitate transport systems involving one or more transportation companies and in which players act as planners or managers. Players confront the same analytic and decision-making problems encountered in practice in real firms, while having access to analytic and decision-making tools incorporated into the simulation games. Lardinois [57] and Lardinois, et al. [58] describe this work, which incorporates techniques of mathematical modelling, operations research and computer graphics within a gaming-simulation approach [59].

Parallel computer technology would provide considerable advantages in the development simulators such as these, for two reasons. The first involves speed: in an environment of training and experimentation, the computer must respond quickly even if the simulation concerns phenomena and decisions that are actually spread out over time. The second is that parallel computing techniques can be used profitably in designing and constructing simulation games in which several individuals participate simultaneously. In a competitive situation, for example, each participant could be assigned to a different processor, so that the various participants in the game would operate concurrently in true competition with one another.

CONCLUSION

Without a doubt, transportation research can and should take advantage of the potential afforded by parallel computing. Several classical base problems in operations research are perfectly suited to this technology; in addition, many well-known solution methods could also benefit from parallel implementations. It is also clear that this change will only occur if computerized tools, adapted to the requirements of parallel computation in operations research, are developed as well. In fact, development in certain advanced areas, such as ITS, depends on a genuine uniting of operations research and computer science. Many of the methodological and computational advances described in this paper have already gone beyond the research stage and should very soon find their way into commercial and industrial applications that meet the new socio-economic and technological requirements of the transportation industry.

REFERENCES

1. Babin, A., Florian, M., James-Lefebvre, L. and Spiess, H. "EMME/2: An interactive graphic method for road and transit planning", *Transportation Research Record*, **866**, T.R.B., National Research Council, Washington, D.C., USA, pp 1–9 (1982).

2. Crainic, T.G., Florian, M., Guélat, J. and Spiess, H. "Strategic planning of freight transportation: STAN, an interactive-graphic system", *Transportation Research Record*, **1282**, T.R.B., National Research Council, Washington, D.C., USA, pp 97–124 (1990).

3. Blais, J.-Y., Rousseau, J.-M. and Lamont, J. "The HASTUS vehicle and manpower scheduling system at the Société de Transport de la Communauté Urbaine de Montréal", *Interfaces*, **20**(1), pp 26–42 (1990).

4. Blewet, W.J. and Hu, T.C. "Tree decomposition algorithm for large networks", *Networks*, **7**, pp 289–296 (1977).

5. Liu, J.W.H. "A graph partitioning algorithm by node separators", *ACM Transactions on Mathematical Software*, **15**, pp 198–219 (1989).

6. Mohr, T. and Pasche, C. "A parallel shortest path algorithm", *Computing*, **40**, pp 281–292 (1988).

7. Pallottino, S. and Scutellà, M.G. "Strongly polynomial auction algorithms for shortest paths", *TR-19/91*, Dipartimento di Informatica, Università di Pisa, Italy (1991).

8. Bertsekas, D.P. "The auction algorithm for shortest paths", *SIAM J. on Optimisation*, **1**, pp 425–447 (1991).

9. Dembo, R.S., Pinar, M.C. and Zenios, S.-A. "A smooth penalty function algorithm for network-structured problems", Report 90-12-05, Decision Sciences Department,

The Wharton School, University of Pennsylvania, PA, USA (1991).

10. Pinar, M.C. and Zenios, S.A. "A comparative study of parallel decompositions for multicommodity flow problems", Report 91-07-07, Decision Sciences Department, The Wharton School, University of Pennsylvania, PA, USA (1991).

11. Gendron, B. and Crainic, T.G. "Relaxations for multicommodity capacitated network design problems", Publication # 965, Centre de Recherche sur les Transports, Montréal, Canada (1994).

12. Zenios, S.A. and Censor, Y. "Massively parallel row-action algorithms in some nonlinear transportation problems", Report 89-09-10, Decision Sciences Department, The Wharton School, University of Pennsylvania, PA, USA (1990).

13. Guélat, J., Florian, M. and Crainic, T.G. "A multimode multiproduct network assignment model for strategic planning of freight flows", *Transportation Science*, **24**, pp 25–39 (1990).

14. Chabini, I. "Des implantations parallèles de l'algorithme d'approximation linéaire pour la résolution du problème d'affectation du trafic", M.Sc. dissertation, Départment d'Informatique et Recherche Opérationnelle, Université de Montréal, Canada (1990).

15. Chabini, I. "Nouvelles méthodes séquentielles et parallèles pour l'optimisation de réseaux à coûts linéaires et convexes", Ph.D. dissertation, Départment d'Informatique et de Recherche Opérationnelle, Université de Montréal, Canada (1994).

16. Florian, M. "An introduction to network models used in transportation planning", in *Transportation Planning Models*, M. Florian, Ed., North-Holland, Amsterdam, pp 137–152 (1984).

17. Pinar, M.C. and Zenios, S.A. "Parallel decomposition of multicommodity network flows using a linear-quadratic penalty algorithm", *ORSA J. on Computing*, **4**, pp 235–249 (1992).

18. Patriksson, M. "Algorithms for urban traffic network equilibria", Thesis no. 263, Department of Mathematics, Institute of Technology, Linköping University, Sweden (1991).

19. Bertsekas, D.P. and Tsitsiklis, J.N. *Parallel and Distributed Computation, Numerical Methods*, Prentice Hall, NJ, USA (1989).

20. Gendron, B. "Implantations parallèles d'un algorithme de séparation et évaluation progressive pour résoudre le problème de localisation avec équilibrage". M.Sc. dissertation, Publication # 761, Centre de Recherche sur les Transports, Université de Montréal, Canada, (1991).

21. Gendron, B. and Crainic, T.G. "Parallel implementations of a branch-and-bound algorithm for multicommodity location with balancing requirements", *Infor.*, **31**, pp 151–165 (1993).

22. Gendron, B. and Crainic, T.G. "Parallel branch-and-bound algorithms: Survey and synthesis", *Operations Research*, **42**, pp 1042–1066 (1994).

23. Crainic, T.G., Gendreau, M. and Dejax, P. "Dynamic and stochastic models for the allocation of empty containers", *Operations Research*, **41**, pp 102–126 (1993).

24. Breton, M. and El Hachem, S. "Decomposition algorithm for stochastic dynamic programs", Publication # G-91-15, GERAD, École des Hautes Études Commerciales, Montréal, Canada (1991).

25. Mulvey, J.M. and Vladimirou, H. "Stochastic network programming for financial planning problems", Report SOR-89-7, Department of Civil Engineering and

Operations Research, Princeton University, NJ, USA (1990).

26. Mulvey, J.M. and Vladimirou, H. "Applying the progressive hedging algorithm to stochastic generalized networks", Report SOR-89-9, Department of Civil Engineering and Operations Research, Princeton University, NJ, USA (1990).

27. Dantzig, G.B. and Glynn, P.W. "Parallel processors for planning under uncertainty", *Annals of Operations Research*, **22**, pp 1–21 (1990).

28. Rumelhard, D. and McClelland, J.L. *Parallel Distributed Processing: Explorations. Microstructure of Cognition*, The MIT Press, Cambridge, MA, USA (1986).

29. Quinton, P. and Van Dongen V. "The mapping of linear recurrence equations on regular arrays", *The J. of VLSI Signal Processing*, pp 95–113 (1989).

30. Drissi-Kaïtouni, O. and Benchekroun Hameda, A. "A dynamic traffic assignment model and a solution algorithm", *Transportation Science*, **26**, pp 119–128 (1992).

31. Kindervater, G.A.P. and Trienekens, H.W.J.M. "Experiments with parallel algorithms for combinatorial problems", *European J. of Operational Research*, **33**, pp 65–81 (1988).

32. Gartner, N.H. and Improta, G., Eds. "Urban traffic networks", *Transportation Research*, special issue, **24B**, pp 407–495 (1990).

33. Sawack, D.J. and Thompson, G.L. "A dynamic space-time network flow model for city traffic congestion", *Transportation Science*, **21**, pp 153–162 (1987).

34. Drissi-Kaïtouni, O. "A model for the dynamic traffic assignment problem", Publication # 702, Centre de Recherche sur le Transports, Université de Montréal, Canada (1990).

35. Drissi-Kaïtouni, O. and Gendreau, M. "A new dynamic traffic assignment model", Publication # 854, Centre de Recherche sur les Transports, Université de Montréal, Canada (1992).

36. Pruul, E.A., Nemhauser, G.L. and Rushmeier, R.A. "Branch-and-bound and parallel computation: A historic note", *Operations Research Letters*, **7**, pp 65–69 (1989).

37. Gendron, B. and Crainic, T.G. "Parallel implementation of bounding procedures for multicommodity capacitated network design", Publication CRT-94-45, Centre de Recherche sur les Transports, Montréal, Canada (1994).

38. Pruul, E.A. "Parallel processing and a branch-and-bound algorithm", M.Sc. dissertation, Cornell University, NY, USA (1975).

39. Trienekens, H.W.J.M. "Parallel branch-and-bound on a MIMD system", Report 8640/A, Econometric Institute, Erasmus University, Rotterdam, Holland (1986).

40. Quinn, M.J. "Analysis and implementation of branch-and-bound algorithms on a hypercube multicomputer", *IEEE Transactions on Computers*, **39**, pp 384–387 (1990).

41. Finkel, R. and Manber, U. "DIB – A distributed implementation of backtracking", *ACM Transactions on Programming Languages and Systems*, **9**, pp 235–256 (1987).

42. Pardalos, P.M. and Rodgers, G.P. "Parallel branch-and-bound algorithms for quadratic zero-one program on the hypercube architecture", *Annals of Operations Research*, **22**, pp 271–292 (1990).

43. Glover, F. "Tabu search, part I", *ORSA J. on Computing*, **1**, pp 190–206 (1989).

44. Glover, F. "Tabu search, part II", *ORSA J. on Computing*, **2**, pp 4–32 (1989).

45. Gendreau, M., Hertz, A. and Laporte, G. "A tabu search heuristic for the vehicle routing problem", *Management Science*, **40**, pp 1276–1290 (1994).

46. Crainic, T.G., Gendreau, M., Soriano, P. and Toulouse, M. "A tabu search procedure for multicommodity location/allocation with balancing requirements", *Annals of Operations Research*, **41**, pp 359–383 (1993).

47. Crainic, T.G., Toulouse, M. and Gendreau, M., "Towards a taxonomy of parallel tabu search algorithms", Publication # 933, Centre de Recherche sur les Transports, Montréal, Canada (1993).

48. Crainic, T.G., Toulouse, M. and Gendreau, M., "Synchronous tabu search parallelization strategies for multicommodity location-allocation with balancing requiremants", O.R. Spectrum, forthcoming.

49. Crainic, T.G., Toulouse, M. and Gendreau, M., "Parallel asynchronous tabu search for multicommodity location-allocation with balancing requirements", *Annals of Operation Research*, forthcoming.

50. Potvin, J.Y. and Rousseau., J.-M. "A parallel route building algorithm for the vehicle routing and scheduling problem with time windows", *European J. of Operational Research*, **66**, pp 331–340 (1993).

51. Foisy, C. and Potvin, J.Y. "Implementing an insertion heuristic for vehicle routing on parallel hardware", *Computers & Operations Research*, **20**, pp 737–745 (1993).

52. Fox, M.S. and Smith, S.F. "ISIS: A knowledge-based system for factory scheduling", *Expert Systems*, **1**, pp 25–49 (1984).

53. Potvin, J.Y. and Rousseau, J.-M. "Constraint-directed search for the advanced-request dial-a-ride problem with service quality constraints", in *Computer Science and Operations Research: New Developments in their Interfaces*, Balci, O. Sharda, R. and Zenios, S.A., Eds., Pergamon Press, pp 457–474 (1992).

54. Potvin, J.-Y., Shen, Y. and Rousseau, J.-M. "Neural networks for automated vehicle dispatching", *Computers & Operations Research*, **19**, pp 267–276 (1992).

55. Ran, B., Boyce, E.D. and LeBlanc, L.J. "Solving an instantaneous dynamics user-optimal traffic assignment model", Urban Transportation Center, University of Illinois at Chicago, IL, USA (1991).

56. Shen, Y. and Potvin, J.-Y. "Simulating the process of multi-attribute choice with neural networks", *Proc. of the IEEE Int. Joint Conf. on Neural Networks*, Baltimore, MD, USA IEEE Press, III-409-414 (1992).

57. Lardinois, C. "Simulation, gaming and training in a competitive multimodal multicompany intercity passenger transportation environment", *J. of the Operational Research Society*, **40**, pp 849–861 (1989).

58. Lardinois, C., Crainic, T.G. and Gendreau, M. "Conception et développement d'un simulateur d'entreprises de camionnage à des fins de formation, d'expérimentation et de transfert technologique", *Les Cahiers Scientifiques du Transport*, **22**, pp 51–64 (1990).

59. Greenblat, C. *Designing Games and Simulation*, Sage Publishing, Beverley Hills, CA, USA (1987).