

Chapter 5

Parallel Computing in Network Optimization

Dimitri Bertsekas

*Laboratory for Information and Decision Systems, Massachusetts Institute of Technology,
Cambridge, MA, U.S.A.*

David Castañón

*Department of Electrical, Computer and Systems Engineering, Boston University, Boston, MA,
U.S.A.*

Jonathan Eckstein

*Mathematical Sciences Research Group, Thinking Machines Corporation, Cambridge, MA,
U.S.A.*

Stavros Zenios

*Decision Sciences Department, Wharton School, University of Pennsylvania, Philadelphia, PA,
U.S.A.*

1. Introduction

Parallel and vector supercomputers are today considered basic research tools for several scientific and engineering disciplines. The novel architectural features of these computers — which differ significantly from the von Neumann model — are influencing the design and implementation of algorithms for numerical computation. Recent developments in large scale optimization take into account the architecture of the computer where the optimization algorithms are likely to be implemented. In the case of network optimization, in particular, we have seen significant progress in the design, analysis, and implementation of algorithms that are particularly well suited for parallel and vector architectures. As a result of these research activities, problems with several millions of variables can be solved routinely on parallel supercomputers. In this chapter, we discuss algorithms for parallel computing in large scale network optimization.

We have chosen to focus on a sub-class of network optimization problems for which parallel algorithms have been designed. In particular, for the most part, we address only *pure* networks (i.e., without arc multipliers). We also avoid discussion on large-scale problems with embedded network structures, like the *multicommodity* network flow problem, or the *stochastic* network problem. Nevertheless, we discuss parallel algorithms for both linear and nonlinear problems, and special attention is given to the assignment problem as well as other problems with bipartite structures (i.e., transportation problems). The problems we have chosen to discuss

usually provide the building blocks for the development of parallel algorithms for the more complex problem structures that we are not addressing. Readers who are interested in a broader view of parallel optimization research — both for network structured problems and mathematical programming in general — should refer to several journal issues focused on parallel optimization which have been published recently on this topic [Mangasarian & Meyer, 1988, 1991; Meyer & Zenios, 1988; Rosen, 1990] or the textbook of Bertsekas & Tsitsiklis [1989].

1.1. Organization of this chapter

The introductory section discusses parallel architectures and broad issues that relate to the implementation and performance evaluation of parallel algorithms. It also defines the network optimization problems that will be discussed in subsequent sections. Section 2 develops the topic of parallel computing for linear network optimization problems and Section 3 deals with nonlinear networks. Concluding remarks, a brief overview of additional work for multicommodity network flows and stochastic network programs, as well as open issues, are addressed in Section 5.

Each of Sections 2–4 is organized in three subsections along the following thread. First, we present general methodological ideas for the design of specific algorithms for each problem class. Here, we present selectively those algorithms that have some potential for parallelism. The methodological development is followed by a subsection of parallelization ideas, i.e., specific ways in which each algorithm can be implemented on a parallel computer. Finally, computational results with the parallel implementation of some of the algorithms that have appeared in the literature are summarized and discussed.

1.2. Parallel architectures

Parallelism in computer systems is not a recent concept. ENIAC — the first large-scale, general-purpose, electronic digital computer built at the University of Pennsylvania — was designed with multiple functional units for adding, multiplying, and so forth. The primary motivation behind this design was to deliver the computing power infeasible with the sequential electronic technology of that time. The shift from diode valves to transistors, integrated circuits, and very large scale integrated circuits (VLSI) rendered parallel designs obsolete and uniprocessor systems were predominant through the late sixties.

The first milestone in the evolution of parallel computers was the Illiac IV project at the University of Illinois in the 1970's. A brief historical note on this project can be found in Desrochers [1987]. The array architecture of the Illiac prompted studies on the design of suitable algorithms for scientific computing. Interestingly, a study of this sort was carried out for linear programming [Pfefferkorn & Tomlin, 1976] — one of the first studies in parallel optimization. The Illiac never went past the stage of the research project, however, and only one machine was ever built.