

SECOND EDITION

**INTELLIGENT  
TRANSPORTATION  
SYSTEMS**

**Smart and Green  
Infrastructure Design**

---

# Mechanical Engineering Series

*Frank Kreith, Series Editor*

---

*Computer Techniques in Vibration*

Edited by Clarence W. de Silva

*Distributed Generation: The Power Paradigm for the New Millennium*

Edited by Anne-Marie Borbely and Jan F. Kreider

*Elastic Waves in Composite Media and Structures: With Applications to Ultrasonic Nondestructive Evaluation*

Subhendu K. Datta and Arvind H. Shah

*Elastoplasticity Theory*

Vlado A. Lubarda

*Energy Audit of Building Systems: An Engineering Approach*

Moncef Krarti

*Energy Conversion*

Edited by D. Yogi Goswami and Frank Kreith

*Energy Management and Conservation Handbook*

Edited by Frank Kreith and D. Yogi Goswami

*The Finite Element Method Using MATLAB, Second Edition*

Young W. Kwon and Hyochoong Bang

*Fluid Power Circuits and Controls: Fundamentals and Applications*

John S. Cundiff

*Fundamentals of Environmental Discharge Modeling*

Lorin R. Davis

*Handbook of Energy Efficiency and Renewable Energy*

Edited by Frank Kreith and D. Yogi Goswami

*Heat Transfer in Single and Multiphase Systems*

Greg F. Naterer

*Heating and Cooling of Buildings: Design for Efficiency, Revised Second Edition*

Jan F. Kreider, Peter S. Curtiss, and Ari Rabl

*Intelligent Transportation Systems: Smart and Green Infrastructure Design, Second Edition*

Sumit Ghosh and Tony S. Lee

*Introduction to Precision Machine Design and Error Assessment*

Edited by Samir Mekid

*Introductory Finite Element Method*

Chandrakant S. Desai and Tribikram Kundu

*Machine Elements: Life and Design*

Boris M. Klebanov, David M. Barlam, and Frederic E. Nystrom

*Mathematical and Physical Modeling of Materials Processing Operations*

Olusegun Johnson Ilegbusi, Manabu Iguchi, and Walter E. Wahnsiedler

*Mechanics of Composite Materials*

Autar K. Kaw

*Mechanics of Fatigue*

Vladimir V. Bolotin

*Mechanism Design: Enumeration of Kinematic Structures According to Function*

Lung-Wen Tsai

*Mechatronic Systems: Devices, Design, Control, Operation and Monitoring*

Edited by Clarence W. de Silva

*The MEMS Handbook*, Second Edition (3 volumes)

Edited by Mohamed Gad-el-Hak

*MEMS: Introduction and Fundamentals*

*MEMS: Applications*

*MEMS: Design and Fabrication*

*Multiphase Flow Handbook*

Edited by Clayton T. Crowe

*Nanotechnology: Understanding Small Systems*

Ben Rogers, Sumita Pennathur, and Jesse Adams

*Nuclear Engineering Handbook*

Edited by Kenneth D. Kok

*Optomechatronics: Fusion of Optical and Mechatronic Engineering*

Hyungsuck Cho

*Practical Inverse Analysis in Engineering*

David M. Trujillo and Henry R. Busby

*Pressure Vessels: Design and Practice*

Somnath Chattopadhyay

*Principles of Solid Mechanics*

Rowland Richards, Jr.

*Thermodynamics for Engineers*

Kau-Fui Vincent Wong

*Vibration Damping, Control, and Design*

Edited by Clarence W. de Silva

*Vibration and Shock Handbook*

Edited by Clarence W. de Silva

*Viscoelastic Solids*

Roderic S. Lakes

CRC Press  
Taylor & Francis Group  
6000 Broken Sound Parkway NW, Suite 300  
Boca Raton, FL 33487-2742

© 2010 by Taylor and Francis Group, LLC  
CRC Press is an imprint of Taylor & Francis Group, an Informa business

No claim to original U.S. Government works

Printed in the United States of America on acid-free paper  
10 9 8 7 6 5 4 3 2 1

International Standard Book Number: 978-1-4398-3518-0 (Hardback)

This book contains information obtained from authentic and highly regarded sources. Reasonable efforts have been made to publish reliable data and information, but the author and publisher cannot assume responsibility for the validity of all materials or the consequences of their use. The authors and publishers have attempted to trace the copyright holders of all material reproduced in this publication and apologize to copyright holders if permission to publish in this form has not been obtained. If any copyright material has not been acknowledged please write and let us know so we may rectify in any future reprint.

Except as permitted under U.S. Copyright Law, no part of this book may be reprinted, reproduced, transmitted, or utilized in any form by any electronic, mechanical, or other means, now known or hereafter invented, including photocopying, microfilming, and recording, or in any information storage or retrieval system, without written permission from the publishers.

For permission to photocopy or use material electronically from this work, please access [www.copyright.com](http://www.copyright.com) (<http://www.copyright.com/>) or contact the Copyright Clearance Center, Inc. (CCC), 222 Rosewood Drive, Danvers, MA 01923, 978-750-8400. CCC is a not-for-profit organization that provides licenses and registration for a variety of users. For organizations that have been granted a photocopy license by the CCC, a separate system of payment has been arranged.

**Trademark Notice:** Product or corporate names may be trademarks or registered trademarks, and are used only for identification and explanation without intent to infringe.

---

#### Library of Congress Cataloging-in-Publication Data

---

Ghosh, Sumit, 1958-

Intelligent transportation systems : smart and green infrastructure design / by Sumit Ghosh and Tony S. Lee. -- 2nd ed.

p. cm. -- (Mechanical engineering series ; 44)

Includes bibliographical references and index.

ISBN 978-1-4398-3518-0

1. Intelligent transportation systems. 2. Green technology. 3.

Transportation--Automation. 4. Transportation--Data processing. 5. Electronics in transportation. 6. Traffic engineering. I. Lee, Tony S., 1970- II. Title.

TA1230.G4565 2010

388.3'12--dc22

2009048023

---

Visit the Taylor & Francis Web site at  
<http://www.taylorandfrancis.com>

and the CRC Press Web site at  
<http://www.crcpress.com>

SECOND EDITION

# INTELLIGENT TRANSPORTATION SYSTEMS

Smart and Green  
Infrastructure Design

Sumit Ghosh  
Tony S. Lee



CRC Press

Taylor & Francis Group

Boca Raton London New York

---

CRC Press is an imprint of the  
Taylor & Francis Group, an **informa** business



# *Dedication*

---

*To my loving family  
[Sumit]*

*To YK for understanding  
[Tony]*





---

# Contents

**List of Figures.....xiii**

**List of Tables.....xvii**

**Preface to the Second Edition .....xix**

**Preface to the First Edition .....xxi**

**About the Authors .....xxv**

**Chapter 1** The State of the Art in ITS ..... 1

- 1.1 The Broad Scope of ITS ..... 1
- 1.2 A Working Definition of ITS ..... 2
- 1.3 Current Status of ITS ..... 2
- 1.4 A Critical Review of the State of the Art in ITS ..... 4

**Chapter 2** New Meta-Level Principles for an Untapped ITS  
Technological Mine ..... 7

- 2.1 Meta-Level Principles..... 7
- 2.2 The ITS Mine and New Mining Implements..... 8
- 2.3 Examples of Formidable Challenges and Amazing  
Opportunities ..... 11

**Chapter 3** Fundamental Issues in Transportation Systems ..... 19

- 3.1 Principal Characteristics of ITS..... 20
- 3.2 Scientific Validation of ITS Designs through Modeling and  
Simulation..... 27

**Chapter 4** DARYN: A Distributed Decision-Making Algorithm for  
Railway Networks ..... 29

- 4.1 Introduction..... 29
- 4.2 The DARYN Approach..... 32
  - 4.2.1 Algorithm ..... 32
  - 4.2.2 Proof of Freedom from Deadlock ..... 35
  - 4.2.3 Modeling DARYN on a Loosely Coupled Parallel  
Processor ..... 36

4.3 Implementation of DARYN on ARMSTRONG..... 38

4.4 Performance of DARYN..... 44

**Chapter 5 RYNSORD: A Novel, Decentralized Algorithm for Railway Networks with Soft Reservation ..... 51**

5.1 Introduction..... 51

5.2 The RYNSORD Approach..... 52

5.3 Modeling RYNSORD on an Accurate, Realistic, and Parallel Processing Test Bed..... 60

5.4 Implementation Issues ..... 66

5.5 Simulation Data and Performance Analysis ..... 68

**Chapter 6 DICAF: A Distributed, Scalable Architecture for IVHS ..... 83**

6.1 Introduction..... 83

6.2 DICAF: A Novel, Distributed, and Scalable Approach to IVHS ..... 90

6.3 Modeling DICAF on an Accurate, Realistic, and Parallel Processing Test Bed ..... 98

6.4 Implementation and Debugging Issues..... 103

6.5 Simulation Results and Performance Analysis of DICAF..... 104

**Chapter 7 Stability of RYNSORD under Perturbations..... 127**

7.1 Introduction..... 127

7.2 Formal Definition of Stability of RYNSORD ..... 130

7.3 Modeling RYNSORD for Stability Analysis..... 133

7.3.1 Implementation Issues ..... 133

7.4 Stability Analysis of RYNSORD..... 133

7.4.1 Error Criteria for Stability Analysis ..... 133

7.4.2 Steady-State Analysis..... 134

7.4.3 Perturbations to the Input Rate and Stability Analysis..... 137

7.4.4 Perturbations to System Characteristics and Stability Analysis..... 138

7.4.4.1 Perturbations to Interstation and Train-to-Station Communications ..... 140

7.4.4.2 Perturbations Relative to the Track Segments..... 147

**Chapter 8 Modeling and Simulation Techniques for ITS Designs ..... 153**

8.1 Introduction..... 153

8.2 Virtual and Physical Process Migration Strategies for ITS Designs ..... 153

- 8.2.1 Virtual Process Migration Strategy ..... 154
- 8.2.2 Physical Process Migration Strategy ..... 156
- 8.3 Software Techniques Underlying the Process Migration Strategies..... 158
  - 8.3.1 Software Techniques Underlying VPM ..... 158
  - 8.3.2 Software Techniques Underlying PPM..... 162
- 8.4 Implementation Issues ..... 165
- 8.5 Simulation Results and Performance Analysis ..... 166
  
- Chapter 9** Future Issues in Intelligent Transportation Systems ..... 173
  
- Chapter 10** Description of the RYNSORD Simulator on CD-ROM and Scope of Experiments..... 175
  - 10.1 Installation ..... 175
  - 10.2 Overview..... 176
  - 10.3 Getting Ready to Run ..... 177
    - 10.3.1 network.out ..... 177
    - 10.3.2 Helper Scripts..... 178
    - 10.3.3 Input Generation ..... 178
    - 10.3.4 Output Files..... 178
    - 10.3.5 Troubleshooting ..... 179
      - 10.3.5.1 How You Know It Is Working ..... 179
      - 10.3.5.2 Problems ..... 180
    - 10.3.6 Track and Communication Failures ..... 180
  - 10.4 Conclusions..... 180
  
- Bibliography ..... 181**



---

# Figures

4.1	Centralized scheduling in traditional approaches. ....	30
4.2	Algorithm for centralized scheduling. ....	31
4.3	The DARYN approach. ....	34
4.4	A potential timing problem in the distributed simulation. ....	38
4.5	The main and initial routines. ....	40
4.6	The read_inputs routine. ....	41
4.7	The train_function routine. ....	42
4.8	The clear_buffer routine. ....	43
4.9	The traveling_trains routine. ....	44
4.10	Four railway networks to evaluate DARYN's performance. ....	45
4.11	Graphs of CPU times versus number of trains: (a) DARYN algorithm, (b) traditional approach. ....	47
4.12	Graph of speedup factor versus number of trains for all four railway networks. ....	48
4.13	Graphs of CPU times for a train versus number of trains for all four railway networks: (a) DARYN algorithm, (b) traditional approach. ....	49
4.14	Graphs of cumulative idle times versus number of trains for all four railway networks: (a) DARYN algorithm, (b) traditional approach. ....	50
5.1	An example railway network in RYNSORD. ....	54
5.2	Computation of primary and secondary paths in RYNSORD: (a) for train Ta, (b) for train Tb. ....	56
5.3	Initiation and progress of reservation processing. ....	59
5.4	Functionality of a station node in pseudocode. ....	61
5.5	Functionality of the special, synchronizing node. ....	62
5.6	A screen shot of the graphical interface. ....	64
5.7	A screen shot of an actual simulation run. ....	65
5.8	A representative railway network: A 50-station subset of the eastern United States railroad network. ....	67
5.9	(a) Average travel time minus ideal travel time averaged over all trains as a function of the lookahead, (b) percentage of trains reaching their destinations within the maximum simulation time. ....	70
5.10	Distribution of actual number of hops for trains in RYNSORD versus approach B versus ideal distribution. ....	71
5.11	Distribution of actual number of hops for trains in RYNSORD versus ideal distribution under (a) low input train density, (b) medium input train density, and (c) high input train density. ....	72
5.12	Average hops minus ideal number of hops averaged over all trains as a function of lookahead. ....	74

5.13 Actual travel time of a train in RYNSORD minus ideal travel time for  
lookahead 4 as a function of assertion time of the train, (a) low input  
traffic density, (b) medium input traffic density, and (c) high input traffic  
density. ....75

5.14 Track utilization for low input traffic density and lookahead 6,  
(a) approach B, (b) RYNSORD. ....76

5.15 Track utilization in RYNSORD for medium input traffic density for  
(a) lookahead 2, (b) lookahead 4. ....78

5.16 Distribution of reservations processed by stations in RYNSORD for  
medium input traffic density as a function of stations for (a) lookahead 2,  
(b) lookahead 4. ....79

5.17 Distribution of computation among the trains in RYNSORD for medium  
input traffic density for (a) lookahead 2, (b) lookahead 4. ....80

5.18 Maximum communications rate (bytes/min) of interstation links. ....82

6.1 IVHS architecture utilizing centralized traffic management center. ....88

6.2 The DICAF architecture. ....95

6.3 Functionality of a DTMC in DICAF. ....97

6.4 Functionality of a vehicle. ....97

6.5 The “main” routine in DICAF. ....101

6.6 The “process message” routine in DICAF. ....101

6.7 “Check wait-to-leave list” function. ....102

6.8 “Update segment” function in DICAF. ....102

6.9 A second representative highway system modeled in DICAF. ....105

6.10 Normal distribution of the desired speeds of vehicles for traffic density  
values of 3, 6, and 8. ....108

6.11 Percentage differential travel times for all vehicles [calculated as (actual  
travel time for each vehicle – ideal travel time)/(ideal travel time) ×  
100%] for (a) density 3, (b) density 6, and (c) density 8. ....110

6.12 Distribution of C.M. for highway segment 7 as a function of the simulation  
time for (a) density 3, (b) density 6, (c) density 8. ....112

6.13 Average C.M. value for all highway segments of DICAF for (a) density 3,  
(b) density 6, (c) density 8. ....114

6.14 Average number of vehicles on the highway segments computed over the  
entire simulation run for (a) density 3, (b) density 6, (c) density 8. ....116

6.15 Percentage differential travel times for all vehicles [calculated as (actual  
travel time for each vehicle – ideal travel time)/(ideal travel time) ×  
100%] as a function of the assertion time of the vehicle for (a) density 3,  
(b) density 6, (c) density 8. ....118

6.16 Desired speeds of vehicles asserted into DICAF for traffic density 8 as a  
function of the simulation time. ....119

6.17 Actual average speeds of vehicles in DICAF for traffic density 8. ....120

6.18 Percentage differential travel times for all vehicles [calculated as  
(actual travel time for each vehicle – ideal travel time)/(ideal travel  
time) × 100%] for (a) { $\mu = 65$  mph,  $\sigma = 10$  mph}, (b) { $\mu = 65$  mph,  
 $\sigma = 5$  mph}. ....121

- 6.19 (a) Number of vehicles connecting to DTMC5 as a function of simulation time, (b) maximum number of vehicles connecting to the DTMCs. .... 122
- 6.20 (a) Number of inter-DTMC flooding messages at DTMC5 as a function of simulation time, (b) maximum inter-DTMC communication messages (send + receive) for every DTMC. .... 124
- 6.21 Modeling and simulation of the 50-DTMC highway system in DICA. Percentage differential travel times [calculated as (actual travel time for each vehicle – ideal travel time) / (ideal travel time) × 100%] for (a) for all vehicles, (b) as a function of the assertion time of the vehicle. .... 125
- 7.1 A representative railway network: A 50-station subset of the eastern United States railroad network. .... 132
- 7.2 Error criterion I for each train as a function of the assertion time for (a) input rate = 0.125, (b) for input rate = 0.14, and (c) for input rate = 0.175. .... 136
- 7.3 Error criterion II as a function of simulation time for steady-state input rates. .... 138
- 7.4 Error criterion I for each train as a function of the assertion time for (a) input rate = 0.05 input rate and perturbation = 0.5, (b) input rate = 0.125 input rate and perturbation = 0.5, and (c) input rate = 0.125 input rate and perturbation = 3.0. .... 139
- 7.5 Error criterion II as a function of simulation time for different input perturbations. .... 141
- 7.6 (a) Error criterion I for each train as a function of the assertion time for the set of eight link failures for 1440 timesteps and input rate = 0.05, (b) error criterion II as a function of simulation time for the set of eight link failures for 1440 timesteps and input rate = 0.05. .... 143
- 7.7 (a) Error criterion I for each train as a function of the assertion time for the set of eight links failed permanently for input rate = 0.05, (b) error criterion II as a function of simulation time for the set of eight links failed permanently for input rate = 0.05. .... 144
- 7.8 (a) Error criterion I for each train as a function of the assertion time for the set of eight link failures for 1440 timesteps and input rate = 0.125, (b) error criterion II as a function of simulation time for the set of eight link failures for 1440 timesteps and input rate = 0.125. .... 145
- 7.9 (a) Error criterion I for each train as a function of the assertion time for the set of eight link failed permanently for input rate = 0.125, (b) error criterion II as a function of simulation time for the set of eight link failed permanently for input rate = 0.125. .... 146
- 7.10 Error criterion II as a function of simulation time for the set of three and set of eight links permanently failed and input traffic rate = 0.05. .... 147
- 7.11 (a) Illustrating communication link failure, (b) illustrating track segment failure. .... 148
- 7.12 Error criterion II as a function of simulation time for communications link and track perturbations for (a) high input traffic rate and finite duration perturbation, (b) low input traffic rate and permanent perturbation. .... 150

- 8.1 Virtual process migration. .... 156
- 8.2 Physical process migration. .... 157
- 8.3 Establishing the static interconnection network during initialization at each node in pseudocode. .... 159
- 8.4 Structure for every mobile entity in VPM. .... 160
- 8.5 A mobile entity remanifested as a process at a node. .... 160
- 8.6 A node intercepts a message encapsulating a mobile entity. .... 161
- 8.7 The function of the scheduler at each node. .... 162
- 8.8 The main program corresponding to each node. .... 162
- 8.9 The migration, disconnection, and reconnection functions of each mobile entity in PPM. .... 163
- 8.10 The main program corresponding to each mobile entity in PPM. .... 164
- 8.11 The main program corresponding to each stationary node in PPM. .... 165
- 8.12 VPM and PPM simulation time as a function of computational load of mobile entities, (a) 10 stationary nodes and 1 data element exchanged at each hop, (b) 5 stationary nodes and 10 data elements exchanged at each hop. The number of stationary nodes is 10, the number of data elements exchanged at each hop is 1, and the mobile unit engages in 1000 connections and disconnections. .... 168
- 8.13 VPM and PPM simulation time as a function of the number of mobile entities for varying computational load values. The number of stationary nodes is 10, the number of data elements exchanged at each hop is 1, and the mobile unit engages in 1000 connections and disconnections. .... 169
- 8.14 VPM and PPM simulation time as a function of number of data elements exchanged at each hop for (a) load = 10,000, (b) load = 100,000. The number of mobile units is 10, and the mobile unit engages in 1000 connections and disconnections. .... 170



---

# Tables

5.1	Input Train Traffic Parameters .....	68
5.2	Comparative Impact of Lookahead on Performance Parameters in RYNSORD under Low Input Traffic Density.....	81
5.3	Comparative Impact of Lookahead on Performance Parameters in RYNSORD under Medium Input Traffic Density.....	81
7.1	Input Traffic Parameters for Steady-State Analysis of RYNSORD.....	135
7.2	Perturbations to Input Rate and System Characteristics.....	138
7.3	Performance Results for Communication Perturbations .....	142
7.4	Performance Results for Track Perturbations .....	149



---

# Preface to the Second Edition

It has been more than 10 years since we sat down to write the first edition, *Intelligent Transportation Systems: New Principles and Architectures* with great excitement, back in 1999. Since then a number of significant and encouraging improvements have been incorporated into transportation systems, worldwide, while a number of areas are continuing to witness rapid deterioration. One undeniable fact stands clear. The promise of intelligent transportation systems (ITS) is greater than ever before, and it is yet to be fulfilled. The congestion, uncertainty, and the lack of a clear direction for the future of transportation continue to persist. Key impediments include the absence of a critical mass of ITS engineer-scientists with the right background and training as well as the lack of adequate research and development thrusts, worldwide. This book is intended to serve as a reminder, within the United States, of not only the development of educational programs focused on ITS directed toward the future transportation workforce but also the professional capacity building initiatives underlying the current transportation workforce. Both of these may be traced to a formidable challenge that we have already begun to encounter in the twenty-first century in virtually every field—the need for a completely new mindset that is intimately connected to interdisciplinary, transdisciplinary, or cross-disciplinary research and education.

Our activities, engagements, and experiences over the past 10 years in ITS and related fields have provided us with a few startling revelations, unique insights, and meta-level principles, which we have incorporated in this revised edition. We had already focused on intelligent decision making and efficiency as the hallmark characteristics of ITS. In this edition, we emphasize the fractal nature of both intelligence and creativity, which serves as a great inspiration toward analyzing the most complex ITS problems at incredible depths and subtlety and promises to yield amazing solutions to nearly impossible problems. Properly cultivated, this fractal attribute offers a viable and practical path toward continuous quantum advancements, ad infinitum. We also observe that the efficiency attribute of ITS is analogous to the ubiquitous “green” color of life on earth in that both ITS and life aim to utilize the least amount of energy and achieve the greatest degree of order, the antithesis of which is entropy or disorder. Our research also reveals that efficiency is equally fractal in nature.

We owe immense gratitude to Dr. Frank Kreith, the series editor and a Fellow of the National Conference of State Legislatures. It was the most pleasant surprise when he kindly called late April 2009 and suggested that we consider revising the book, given that there is significantly greater interest in ITS today than back in 1999. There is an obvious explanation for the growing interest, namely, the billions of dollars allocated by U.S. Congress to repair roads, bridges, and infrastructure under the American Recovery and Reinvestment Act of 2009. There is, however, a deeper unseen reason. The United States has a rich history of creativity, innovation, and entrepreneurial commercialization of ideas. Of even greater importance, however, is that this nation has a history of breaking out of economically distressed times

with incredible explosions of efficiency and innovation. Between 1930 and 1950, radar, nuclear energy, interstate highways, transistors, and computers had reshaped the economic landscape of the United States as well as that of the entire world. In 2010, the discipline of ITS with its unlimited scope stands ready to serve in this great tradition.

In this edition, we have included two new chapters. Chapter 1 critically reviews important developments and implementations of ITS, worldwide, and presents key challenges continuing to confront ITS. Chapter 2 is organized into three parts. The first part presents the meta-level principles and insights, while the second part explains how the insight and meta-principles may be applied to real-world ITS problems. The third part poses a number of formidable challenges for ITS, which, if and when solved utilizing the meta-level and ITS principles, will imply a quantum leap in technology and unlimited societal benefits. It is our hope that transportation personnel in municipalities, counties, and states across the United States as well as in India, China, and other nations of the world will review the ideas presented in this book and seriously consider adopting them in their current and planned roads, highways, bridges, and infrastructure projects. Among the promised benefits, the quality and operational lives of these projects may be extended far into the distant future. The fundamental characteristics and principles of ITS systems in Chapter 3 have been further elaborated in this revised edition.

As an industry, ITS holds tremendous potential for private–public partnership. However, it is yet to be studied thoroughly as an industry. ITS is vast and its many dimensions include technical, operational, policy, privacy, standards, business, financing, political, and how it can be linked with other industries to foster growth and wealth creation. This book focuses primarily on the technical and operational issues and, secondarily, on the policy issues of ITS.

We are deeply indebted and grateful to Dr. Tim Lomax of the Texas Transportation Institute at Texas A&M University and Prof. Chelsea White of Georgia Tech for their insightful, constructive, and candid reviews of the book.

We are also deeply grateful to Mr. Raj Ghaman, former manager of research in the ITS Joint Program Office at the U.S. Department of Transportation, Washington, D.C.

Finally, our sincerest appreciation to Joe Clements and the entire editorial and production staff at CRC Press for their continued and enthusiastic support.

**Sumit Ghosh**  
**Tony S. Lee**

---

# Preface to the First Edition

The notion of transportation refers to the movement of people and goods across nontrivial geographical distances, and its history is as old as our civilization. Its scope is enormous, ranging from people walking on the earth's surface to carts and chariots driven by animals, automobiles, trains, airplanes, and ships. Despite the great diversity in the modes of transportation, a unique characteristic emerges—the constant effort that has continued throughout history to improve its efficiency by providing information and guidance to the constituent entities. In the past, information was carried by people and by material in the form of messages and, as a result, the rate of propagation of information was closely related to that of the mode of the transportation. With the emergence of electromagnetic communication, in the last century, the discipline of transportation experienced a remarkable transformation. Information and guidance could now be provided much faster than the actual rate of movement of the people and goods, implying a qualitative improvement in the transportation system. Thus, in the railway system, it became a standard practice for a station master of a station A to “wire” ahead to the station master of station B the impending arrival of the train that has just passed station A. For many reasons, including simplicity and the desire to maintain consistent control, the information and guidance providers of the transportation systems evolved as centralized units. A central control would gather information about every entity within a specific transportation system and provide guidance and information to them as necessary. When systems were isolated, small, and few in number, decentralized control was the norm. Ironically, as systems grew larger due to consolidation, etc., control became centralized.

Toward the end of the twentieth century, the number of the constituent units in many transportation systems, especially the highway system, have simply skyrocketed. In the United States alone, there are more than 200 million vehicles and more than 4 million miles of paved roads. In addition, during specific times such as “rush-hour,” the utilization of certain subsets of many transportation systems frequently exceeds the capacity, leading to gridlock and related problems. Furthermore, the average person today is far more demanding when it comes to his/her freedom of choice and flexibility. For many transportation systems, the cost of expanding the existing infrastructure is prohibitively expensive. The result is a major strategic shift from building more infrastructure to providing timely information and high-quality guidance toward improving the quality of transportation within the confines of the existing infrastructure. That being said, some infrastructure building, where absolutely necessary, must be continued, subject to greater thoughtfulness and guided scientifically through modeling and simulation. In turn, the centralized paradigm is confronted with overwhelming challenges, both from the sheer number of entities, the demand for faster and accurate information, and the need for high-quality guidance.

To serve the needs of the future, beyond the year 2000, transportation systems must necessarily undergo another radical transformation, namely from the centralized

paradigm to the asynchronous, distributed paradigm, which will integrate fast computers and high-performance computer networks through novel computer algorithms. This book focuses on the fundamental principles in future transportation systems that, in turn, give rise to new system architectures. The book argues strongly the need to design innovative and creative approaches to transportation problems, utilizing the fundamental principles, followed by thorough scientific validation. It stresses both the need for computer modeling and simulation of a representative system to help design and validate such complex, large-scale systems and the design of new performance metrics to estimate the performance of these systems. Although the scope of transportation systems is very broad, this book focuses in detail on two ubiquitous transportation modes—highway and railway systems. It presents the basic principles pertaining to travel-related decision making that includes coordination, control, and routing. These principles constitute the core of and apply equally to all transportation systems, including passenger air travel, air freight, personalized rapid transit, and so on. While these principles constitute one of several key advances necessary to usher transportation systems into the next century, other relevant issues such as driver behavior analysis, human factors, congestion originating in human behavior, and other traffic factors are beyond the scope of this book.

The development of this book has been motivated by two reasons. The first is my frequent frustration with the current transportation systems and a sincere belief that the technology that is needed to alleviate the problem is already here. Our frustrations, which we are certain we share with millions of our fellow human beings, include traffic congestion, lost luggage during air travel, misinformation about train arrivals at stations and public bus arrivals at stops, and driving in an unknown city late at night and getting lost. The General Electric locomotive building unit reports, according to Carley in the *Wall Street Journal*, dated June 29, 1998, that locomotives sit idle as much as 40% due to bottlenecks in the rail corridors that stem from poor information, coordination, and control. We are convinced, through our fundamental analysis of transportation systems and our own research, that novel distributed control algorithms constitute the most logical choice for future transportation system architecture design. In our vision, we see a significant and increasing role for novel computer control algorithms and large-scale, distributed software in future transportation systems, providing an entirely new range of personalized, travel-related services, qualitatively enhancing the efficiency of movement of each entity throughout the system, sophisticated control of ramp meters at highway entrances, coordinating traffic flow in highways and street surfaces, providing personalized rapid transit services, and fostering a safe environment.

Second, we sincerely believe that to realize a qualitative jump toward intelligent transportation systems, the future highway engineer and traffic specialist must receive interdisciplinary training in civil and electrical engineering, transportation engineering and planning, human factors, and, most importantly, computer science and engineering (CSE). Basic knowledge in the areas of distributed systems, algorithm design, networking, computer modeling, and distributed simulation within CSE are especially important, for they hold the key to understanding the most complex,

top-level system architecture. In the January 1998 issue of *Traffic Technology International*, while Prof. John Collura of the University of Massachusetts presented some details on a new undergraduate and graduate curricula in response to the recent ITS movement, Prof. Chelsea White of the University of Michigan noted that the nature of future transportation problems is essentially interdisciplinary. Recently, Arizona State University has launched a Graduate Interdisciplinary Certificate program in transportation systems, supported by faculty from planning and landscape architecture, civil and environmental engineering, geography, and aeronautical management and technology. The ideas and principles in this book also underlie the Autonomous Decentralized Transport Operation System (ATOS) that controls the world's largest transportation system—East Japan Railway Company. Developed by Hitachi, ATOS includes 5000 autonomous computers that control more than 6200 trains/day [1].

Chapter 3 examines the essential nature of all transportation systems, especially from the perspective of future needs, and presents the fundamental principles that emerge from the analysis. This chapter presents key design issues of future systems including the control algorithms, the nature of the interactions between the different constituent entities of the system, and the network that interconnects the entities. It also underscores the role of modeling and simulation in the design of future systems. Chapters 4 through 6 present a number of case studies that encapsulate the principles outlined in Chapter 3. Chapter 4 describes in detail the first study of a novel, distributed approach to routing of trains in a railway network. This chapter begins with an explanation of the traditional approaches to train control and presents a detailed, critical review of the literature. It then presents a new, distributed control algorithm, DARYN. The algorithm is first modeled for a small-scale railway network and then simulated on Armstrong, a loosely coupled parallel processor to yield performance measures. Chapter 5 presents a highly sophisticated algorithm, RYNSORD, for efficient scheduling and congestion mitigation in railway networks. In RYNSORD, every train utilizes lookahead to dynamically replan its route and, at every stage, reserves one or more tracks prior to utilizing them. The reservation process is characterized as “soft,” that is, less abrupt and more flexible in negotiation between the trains and stations, in contrast to the traditional, rigid, “hard” reservations. RYNSORD is modeled for a subset of the Eastern United States railroad network and simulated on a network of 65+ SUN Sparc 10 workstations. In this chapter, analysis of the performance data focuses on a comparison of the train travel times and a unique metric in the literature—quality of the routing decisions. Chapter 6 introduces DICAF, a novel architecture for IVHS wherein the overall task of dynamic route guidance and intelligent congestion mitigation is distributed among every entity, that is, automobiles and highway infrastructure, of the IVHS system. The term IVHS is utilized in the limited context of vehicles in highways while the term ITS commands a broader scope and encompasses all transportation systems. DICAF utilizes a continuous function—congestion measure, to influence route guidance, and presents a novel metric that contrasts DICAF's performance, for a realistic highway system, against the absolute best. While DICAF and RYNSORD share the common goal of achieving efficient allocation of system resources, they employ different strategies, which, in turn, stem from their basic dif-

ferences. While trains are confined along tracks, which must be reserved exclusively prior to use, automobiles enjoy greater flexibility in sharing the road with each other. Chapter 7 presents a systematic and detailed study of the stability of the novel computer algorithms for transportation in the presence of perturbations. This analysis is key to understanding the robustness and resilience of the complex transportation systems. It utilizes RYNSORD to illustrate the basic principles. Chapter 8 presents new techniques that have been developed for the modeling and simulation of ITS designs.

Thus, the organizational philosophy of this book is as follows. While Chapter 3 presents the general principles, Chapters 4 through 6 address actual problems that serve as concrete examples from the real world. The thinking is that, despite the basic underlying similarity, each transportation problem is unique in its nature, characteristics, and requirements, and needs a custom solution to be developed that is subject to the general principles. It is hoped that the reader will internalize these general and particular concepts and synthesize innovative solutions to future problems.

This book is intended for graduate students in transportation engineering and computer science; researchers in the different disciplines of transportation; practitioners in railways, highway systems, and aviation; and policy makers for transportation infrastructure. It may be used as one of a limited number of textbooks in a graduate course in advanced transportation engineering. In addition to the theoretical concepts, computer modeling, and experimental analysis, this book provides to the readers executable simulators, included in the accompanying CD-ROM, for each of the systems described in the Chapters 5 through 7. The aim is to enable the reader execute the simulations on a network of Linux workstations for different choices of parameters, network configurations, and input traffic data to gain a deeper understanding through hands-on experimentation and experience. Instructors may also use the simulators as a part of a laboratory environment for students to work on laboratory exercises. Chapter 10 describes the use of the simulators. While these simulators are meant for academic use and are limited in scope, the reader may contact [sumit.ghosh@ieee.org](mailto:sumit.ghosh@ieee.org) for pointers to industrial grade versions of the simulators.

The authors are indebted to many individuals who have contributed greatly to their understanding of the transportation systems. They include the first author's former students, Raj Iyer, M.D., Noppanunt Utamapathei, and Kwun Han who had pursued their degrees with him during his tenure at Brown University; Rick Backlund of the Federal Highway Administration, Boston; and Robert Shauver of the Rhode Island Department of Transportation. To the many authors of the papers in the transportation literature that helped us understand the discipline, we express our sincere gratitude. We are also grateful to the funding agencies of the U.S. Department of Defense for their continued support. Last, we express our sincere thanks to the staff at CRC Press for their enthusiasm and meticulous efforts.



---

# About the Authors

**Sumit Ghosh** received his BTech degree (1980) in electrical engineering from the Indian Institute of Technology at Kanpur, India, and his MS and PhD (1984) degrees from the Computer Systems Laboratory (joint EE and CS departments) at Stanford University, California. Sumit's industrial research training has been at Bell Labs Research (Holmdel, New Jersey), Fairchild Advanced Research and Development Labs (Palo Alto, California), and Silvar-Lisco (Menlo Park, California). He has held academic positions at Brown University, Arizona State University, and Stevens Institute of Technology, and is currently at University of Texas at Tyler. He is the recipient of the 2004 IEEE Computer Society's Technical Achievement Award, a Fellow of the Society for Design and Process Science (SDPS), and he has been elected in 2008 to serve as president of SDPS, the leading society for transdisciplinary research and education. He is the primary author of five reference books: *Hardware Description Languages: Concepts and Principles* (IEEE Press, 2000); *Modeling and Asynchronous Distributed Simulation of Complex Systems* (IEEE Press, 2000); *Intelligent Transportation Systems: New Principles and Architectures* (CRC Press, 2000); *Principles of Secure Network Systems Design* (Springer-Verlag, 2002); and *Algorithm Design for Networked Information Technology Systems: Principles and Applications* (Springer-Verlag, 2003). He is the primary coeditor of the book, *Guarding Your Business: A Management Approach to Security* (Kluwer Academic Publishers, 2004).

**Tony S. Lee** received the combined ScB and AB degree from Brown University, Providence, Rhode Island, in computer engineering and political science, and the ScM and PhD degrees in electrical engineering also from Brown University. He is presently a leading scientist at Yahoo Corporation. Tony has published more than a dozen peer-reviewed journal and conference papers in fields ranging from transportation systems, banking systems, and military command and control. He is also a coauthor with Sumit Ghosh on *Modeling and Asynchronous Distributed Simulation: Analyzing Complex Systems* (IEEE Press, June 2000). Prior to his current position, Tony served as senior staff engineer in the Data Transformation Group at Vitria Technology in Sunnyvale, California, working in the area of distributed software development and Internet technologies. Prior to his tenure at Vitria Technology, Tony held a research position with NASA at Ames Research Center, Moffett Field, California. At NASA, he was involved with several high-speed network research projects including Network Quality of Service and the Information Power Grid. Tony is also the author of the ATMSIM software package, a high-fidelity, distributed network simulator that enables network researchers to develop realistic models of new protocols and algorithms and identify complex interactions.

---

# 1 The State of the Art in ITS

## 1.1 THE BROAD SCOPE OF ITS

Since 1999, it has become increasingly clear that the issue of efficiently moving people and goods is far more complex than previously imagined. The problem is not just confined to surface transportation, that is, vehicles and roadways; it affects trains, passenger planes, air cargo, ferries, ships, and all available and currently utilized mode of transportation. This realization had helped foster the broader notion of intelligent transportation system (ITS), superseding intelligent vehicle highway systems (IVHS). Two researchers, Kan Chen and Bob Ervin, conceived the acronym IVHS and were instrumental in expanding its scope and renaming it ITS (Prof. Chelsea White, pers. comm., 2010). Clearly, to synthesize a genuine and practical solution at the national level, possibly international level, one must adopt a holistic approach that takes into consideration complex and asynchronous interdependencies between the many transportation modes and guided by a fundamental goal, namely, minimize the transit time for all travelers and merchandise in transit, subject to fair distribution of the available resources. In the not-too-distant future, space travel to the moon and other artificial satellites and planets in our solar system will become routine and may need to be accommodated within the ITS system. Their unique attributes should be taken into consideration today, while planning out the fundamentals of the ITS architecture, so that integration will be seamless in the future.

Underlying the Intermodal Surface Transport Efficiency Act (ISTEA); the Transportation Equity Act for the 21st Century (TEA-21); the Safe, Accountable, Flexible, Efficient Transportation Equity Act: A Legacy for Users (SAFETEA-LU); and the ITS America movements is the general belief that the solution to this complex problem rests on two key scientific and engineering advances. The first is the increased availability of computing power in the form of powerful desktop workstations and mobile laptops, palmtops, and handheld personal digital assistants (PDAs). The second is the increasing availability of networking, both wire-line and wireless, for communications and control. However, the mere availability of computing power and networking does not automatically guarantee a solution. For instance, Carley [2] from the General Electric locomotive-building unit reports that locomotives sit idle as much as 40% due to bottlenecks in the rail corridors stemming from poor information, coordination, and control. With annual fuel budgets in the late 1990s for many of the railroad companies running in excess of \$800 million, and the need to leave diesel locomotive engines running even when idle, the extent of fuel wastage is staggering. By any standards, such an inefficiency is not sustainable. The key to successful solutions to complex transportation problems, presented in this book, stems from a comprehensive understanding of the control and coordination algorithms. While abstract, these

algorithms serve to unify the computing and networking resources in a synergistic, nontangible manner.

## 1.2 A WORKING DEFINITION OF ITS

ITS is an evolving scientific and engineering discipline whose primary goal is to minimize the travel time of all travelers and merchandise while ensuring safety, through fair distribution of available resources, especially under the scenario of increasing travel speeds, a significantly large number of travelers, and a high demand for precise and timely information by travelers. To achieve its goal, ITS must bring about a seamless and natural integration of the different modes of transportation, including vehicular traffic, trains, cargo air transport, passenger air transport, marine ferries, and others through asynchronous distributed control and coordination algorithms, subject to social norms, policies, and guidelines. As a result of the integration, the traveler will (1) gain access to accurate status information of any transportation mode from any point in the system, (2) compute the most efficient route or reroute across all different transportation modes by processing the available information through personalized decision aids, and (3) be permitted to effect reservations, dynamically, even while en route, on any transportation system. ITS encompasses the subareas of transportation management, which subsumes interstate roadways management and traffic signaling; travel management, which includes multimodal traveler information; public transportation and transit management; safety management, which subsumes incidents, railroad grade crossings, and emergency services; advanced vehicle control; and fare payment and toll collection.

## 1.3 CURRENT STATUS OF ITS

Despite the ISTEA act of 1991, to virtually all drivers today, the current interstate and state highway system continues to be a source of great frustration, the principal reason being congestion. The problem is not new; congestion has been around for a very long time. In London, the average vehicular speed in 2003 was 17 mph, unchanged from 100 years ago when horse-drawn carriages were around. In Manhattan, pedestrians today average a speed of 3.3 mph, which ironically is higher than that of vehicular traffic at 3.1 mph. Based on a ten-year study of the five most congested areas, Los Angeles, Washington, D.C., San Francisco–Oakland, Miami, and Chicago, Lomax reported [3] in 2002 that the average travel delay, in the form of idle waiting, equates to one workweek per year per driver and the cost of wasted fuel from idling totals \$8.6 billion annually in Los Angeles alone. The worst travel delay equals 75 hours per year in the San Bernardino–Riverside area while the highest per capita fuel wastage of \$860 occurs in Washington, D.C. The secondary effect of increased congestion is that the U.S. interstate highway system and major roads in the United States are crumbling [4]. Disintegrated sections of the pavements are not only leaving behind potholes that are dangerous to vehicles at high speeds, but the debris is likely to turn into flying missiles, causing accidents and incidents. To avoid congestion, heavy-duty tractor-trailers in excess of 100,000 pounds are choosing to use rural roads more and more, which, in turn, are falling apart sooner since they were only designed for light traffic, under 10,000 pounds.

Progress in ITS may be viewed from two distinct perspectives. The first perspective, a review of the scientific literature, reveals research in the areas of system planning, vehicular traffic modeling, system evaluation, vehicle tracking, autonomous driving and GPS-based guidance, signal control, braking, lane detection and steering control, intelligent cruise control, disseminating road work information to drivers, noise pollution, platooning, simulators for evaluating ITS systems, and training ITS personnel. While copper cables are increasingly being replaced with wireless and optical fibers to improve communications in rapid transit systems, the duration of the red, amber, and green signals in traffic lights are constantly being recomputed and adjusted to enhance traffic flow and pack vehicles more closely on the roadway. Under autonomous maneuvering studies for lane following and lane changing scenarios, test results reveal that linear cameras at 1000 Hz are highly effective, providing relative distance accurate down to 1 mm at a distance of 10 meters, speed of 60 km/h, and deceleration of 2 m/s<sup>2</sup>. Perhaps the most important accomplishment in ITS to date has been in the area of providing consumers with relevant information, under the rubric of information technology. The availability of arrival and departure information of flights and trains through the Internet from homes, businesses, and on mobile phones, worldwide; automatic notifications of delays and cancellations; facility to reserve seats and call taxicabs through SMS on mobile phones; and the ability to get street directions and maps from a starting point to a destination station on the Internet represents remarkable privileges.

The second perspective consists of a critical examination of the products and services that have been introduced into the market by the ITS industry since 1991. For example, in many U.S. cities, there are toll-free phone numbers (e.g., Boston) and public web sites (e.g., Seattle) where one can obtain Urban Mobility Report (Dr. Tim Lomax, pers. comm., 2010) with up-to-date information on the congestion states of the roadways. These services are immensely valuable but there are three key problems. First, the information is limited to the major roadways, which are generally few in number, while that for the many more secondary roads is either missing or outdated. Second, the time necessary to broadcast complete information on all of the highways is usually so long that its usefulness to any given driver is severely limiting. Third, the information is not dynamic in that it may change substantially due to accidents or incidents while the driver is in transit. Today, a number of trucks, airport shuttle service vans, and luxury automobiles are equipped with GPS-based navigation aids. The principal difficulty again is the lack of accurate and dynamic information regarding congestion and road closures in a timely manner. A number of highway authorities have turned to AM radio to broadcast ramp closures, lane constrictions due to construction, and other relevant highway conditions. This effort, while of great value, has encountered problems where drivers are unable to listen to the broadcast in a timely manner due to electromagnetic interference. Among many transportation communities, there is a strong drive to mount cameras along highways and freeways, feed the signals back to a centralized traffic control center, and monitor congestion. While installation and maintenance costs are high, the cameras can be immensely valuable in zooming onto accident sites and guiding en route police and paramedics or reading license plates of vehicles under suspicion. However, relative to the issue of congestion monitoring and control, other devices including optical

fibers embedded underneath the pavements are far superior in that they are relatively inexpensive and highly reliable. There is a growing private sector travel speed data market. Electronic toll collection systems such as the E-ZPass in the eastern United States and along the Interstate 95 corridor have been viewed by many in the ITS community as a major achievement. Delays are reduced significantly as vehicles whiz by at speeds of 15 mph, avoiding the need to come to a complete stop to conduct cash transactions. In reality, however, E-ZPass's benefits are realized only when roadways are uncongested. Under realistic road conditions, its performance has been, at best, mixed. On any given weekday in some parts of the country, especially during rush hour, thousands of cars hurl toward the toll booths, dangerously zigzagging their way from their current lanes toward the appropriate toll lanes, some slowing down, others coming to a dead stop, in effect creating a highly stressful and unsafe environment. While the queues at the cash toll booths are significantly long, E-ZPass toll lanes barely register average vehicle speeds of more than 2 or 3 mph. Beyond the toll booths, the combination of the abrupt reduction of 10 toll lanes down to 3 or 4 highway lanes plus drivers accelerating quickly trying to make up for lost time creates an even more dangerous environment. The concept of toll collection, electronic or otherwise, as currently deployed, is inefficient in that it slows down traffic unnecessarily, is unsafe in that it fosters a dangerous driving environment, and is environmentally unsound, all of which fundamentally contradicts the basic goals of ISTEA and the objectives of ITS.

Our research into railroad algorithms had revealed, as early as 1992, that the idle-waiting times of trains under centralized scheduling is a factor, between 5 and 40 times, higher than that under distributed scheduling. In 1998, Carley had written that as high as 40% of the locomotives stand idle at the siding stations, stemming from lack of instructions. Yet, as recently as January 2010, we learned from a leading US DOT official that railroads continue to rely on centralized control. For example, throughout the eastern seaboard, CSX trains are centrally controlled from Jacksonville, Florida, and, at any given time, long tracks, stretching 20 miles in the rural parts of the Carolinas, contain only a single train. The result is inevitable congestion and delay.

## **1.4 A CRITICAL REVIEW OF THE STATE OF THE ART IN ITS**

Extensive, engaging, and candid discussions with leading officials from the U.S. Department of Transportation (DOT), the U.S. Federal Highway Administration, and the state-level departments of transportation, followed by subsequent careful analysis have revealed that the lack of tangible progress in ITS may be traced to a handful of formidable and persistent challenges. It is very important that policy makers and transportation personnel discuss, debate, and consider these continuing issues. Otherwise, progress in ITS may remain stalled, which may eventually cause the public to become disillusioned and abandon ITS prematurely.

The first challenge is that there is a strong inherent barrier to envision how ITS can bring to fruition revolutionary changes that will yield quantum improvement to our quality of life and exceptional benefits to society today and well into the future. History teaches us that the single greatest obstacle to our progress is our own inability to envision the future and our tacit refusal to believe in the promises. Where belief

is lacking, there is little motivation to expend the enormous effort that is required for success. As a result, even the greatest idea will fizzle. One fact that may help us immensely in our struggle is as follows: great teachers and thinkers, across nations, cultures, and time have always stressed the importance of imagination. Today, most of the mainstream scientists and engineers dismiss it as a soft science, in contrast to the hard sciences, namely, physics, mathematics, and so on. Today, imagination has no place in any formal higher educational curriculum anywhere in the world. Ironically, however, every great concept and principle, without exception, in any field including science, mathematics, and engineering has begun its journey in the world of intuition and imagination. It is only through great effort that these concepts had emerged to become the foundation and pillars of society. The wheel of the chariot; weapons that could be hurled at a distance; the entire field of number systems; grammar in natural languages; rhythms in music; Fourier analysis; phasor analysis, which has been one of the most difficult concepts in electrical engineering; transistors; discrete digital computers; the Dolby noise reduction system; dynamite; and even the mighty U.S. dollar, which rivaled and defied gold to become the world's first fiat reserve currency, all started out their lives as figments of imagination. This phenomena has been the source of the phrase, "humble beginnings." Can any one of us imagine the world today without these invaluable concepts and the products that they had helped create? In Chapter 2, we will make a strong effort to elaborate on this phenomena and explain how we can embrace it. FIAT is a Latin term, implying, let it be done. FIAT money is currency whose value is derived solely from the corresponding government's decree that the money is legal tender. It bears no intrinsic value; it is not based on any tangible commodities, including gold, silver, and precious metals; and confidence in it stems from the faith and trust in the nation's economy and stability.

The second challenge may be characterized as a self-fulfilling prophecy. Leading transportation personnel have come to believe that the real-world problems are so complex that they can never be solved by ITS. Many transportation officials have shared their frustrations with us in different contexts and flavors. A representative remark came to us a few years earlier, when one of the top DOT officials made the following observation, which is paraphrased here. It is not possible to obtain realistic traffic demands for any ITS project; thus, the knowledge of the necessary design decisions and options is elusive and, therefore, the pursuit is pointless. Consider the following real anecdote. In 1995–1996, the section of Interstate 10 (I-10) from the Phoenix airport to Queen Creek in Arizona had only two lanes each way and was often severely clogged during rush hours. Arizona was experiencing phenomenal growth and new communities were popping up everywhere along the section of I-10. At a great expense and inconvenience to all, the Arizona Department of Transportation embarked on an ambitious project to widen I-10 to three lanes each way. After approximately one year of work, the freeway was finally reopened with all six lanes. At first, it was an amazing improvement. Congestion had virtually disappeared and commute along I-10 became a breeze. Within six months, congestion had returned with a vengeance and the situation was back to square one. The population explosion in barely a year and a half had brought the six-lane I-10 to a grinding halt. Thus, the frustrations of the transportation officials are clearly understandable. Fortunately,

however, we had introduced a new principle [5], the indirect determination of input stimuli; or the stability criterion, to mitigate this precise problem. It is based on sound engineering and logical principles and reflects new and practical thinking. The principle is explained in this book.

The third challenge relates to education. In late 2000, a top research official from the DOT explained to us that the DOT urgently requires sophisticated simulators to study and assess ITS projects. They had commissioned a leading research institute to evaluate the 51 top-rated ITS simulators that were available from top commercial vendors and one leading academic institution. Following extensive testing, both the institute and the DOT concluded that none of the simulators were acceptable. This incident vindicated our long-held professional opinion, namely, a thorough simulation study of any project requires a team of individuals with genuine knowledge and expertise to develop a simulator from scratch, tailored to the needs of the project. No “canned” simulator can be expected to offer the level of flexibility necessary for an ITS project. By definition, the full dimension of a complex problem may never be known a priori, implying that it is very difficult, if not impossible, for a “precooked” simulator with a limited set of options to anticipate all of the sophisticated needs of a given problem. The U.S. Department of Defense, especially the U.S. Army, has spent hundreds of millions of dollars, if not billions, over the past 30 years to perfect canned software simulators that may be used in both the planning and operational phases of the battlefield. Candid discussions with leading military officials reveal that progress can be, at best, described as sputtering. It is our belief that if modeling and simulation expertise had been prevalent in the military at all levels, operational plans would be thoroughly studied prior to engaging the enemy and the many occurrences of mission failures and painful fratricides could have been averted. The DOT’s finding relative to the precooked simulators strengthens our thinking; namely, in order to maintain leadership in ITS research and high-tech industry, the United States needs a serious and extraordinary initiative to (1) bring behavior modeling and asynchronous, distributed simulation into mainstream research and education, and (2) cultivate real talent in this field through training students at the BS through PhD levels and beyond.

The fourth challenge is manifest through the absence of proper nurturing of ITS research and development and, worse, the unwitting imposition of self-defeating incentives and priorities on transportation personnel. Since 1997, enormous sums of money and effort have been poured into installing cameras at traffic light intersections, for the purpose of monitoring congestions and detecting accidents. In reality, automated devices and mechanisms to monitor congestion have been available for a long time that cost, relative to the deployment of the overhead cameras, only 1% to install and whose operational cost is less than 0.01% over the life of the system. While these facts are not unknown to transportation officials, the key reason for such expense is that a false sense of accomplishment and progress is created. While cameras by themselves do not represent true ITS advancement, their obsessive nationwide deployment definitely serves to distract us from the real and necessary improvements. In this era of extreme competition for high-tech industrial dominance, we must provide true freedom to ITS engineer-scientists so that they can diligently and patiently pursue the real and formidable challenges in ITS.



---

# 2 New Meta-Level Principles for an Untapped ITS Technological Mine

## 2.1 META-LEVEL PRINCIPLES

Over the past 10 years, we have acquired a unique insight and have cultivated a number of new meta-level principles, which we present here. The insight stems from an observation that virtually all great inventions and discoveries, exceptions notwithstanding, have been serendipitous in nature, where the inventor or discoverer had accidentally stumbled upon it, as if through pure chance. Examples include the discovery of penicillin, invention of the transistor, dynamite, and others. We also observe that great inventions and discoveries have been few and far in between and that we have merely played a passive role in the evolution of nature. Today, we stand at a unique crossroad in the history of human civilization. For the first time in history, we have the wherewithal to enable deliberate discoveries and inventions. All we need is the will. The enabler we refer to here is the computing engine, of which today's digital discrete computer is one of the best known manifestation.

To understand how computers will help vindicate our claim, we will examine the fundamental nature of the computing engine that yields the following set of meta-level principles. First, the computing engine is, in essence, an amplifier of the human mind, with the amplification factor at a billion (and increasing). The human mind is characterized by two fundamental, high-level functions, namely, the ability to make decisions and to communicate. The computing engine embodies both of these characteristics, except that it can execute them very, very quickly. Second, any idea that we can imagine, any idea at all, may now be implemented in the abstract world of the computer and studied in great detail through behavior modeling and asynchronous distributed simulation. It does not even matter if the idea contradicts the known laws of the physical and social world. It can still be realized in the abstract world of the computer, a totally new dimension that we refer to as cyberspace. Since modeling and simulation can execute extremely fast, we can examine, explore, and study the most complex ideas in our own lifetime and bring them to fruition in the real world. Third, instead of patiently waiting for nature to confront us with problems such as SARS, AIDS, or a super volcanic eruption, and then trying to find a solution, reactively, we can imagine the kinds of challenges we are likely to face in the future and proactively start on our journey to find the right solutions. We will have time to assess the quality of the solutions, refine them, and explore the unintended consequences, if any.



Prof. C.V. Ramamoorthy, Emeritus Professor at University of California, Berkeley, believes that we need to create an entirely new subdiscipline in our scientific and engineering curriculum and label it “Needs Engineering.” He believes, and we concur, that properly conceived conveniences of today will evolve into tomorrow’s societal needs. The fourth meta-level principle is that success can only be achieved if we seriously and passionately embark on a whole new paradigm—transdisciplinary research and education.

## 2.2 THE ITS MINE AND NEW MINING IMPLEMENTS

Although the meta-level principles, introduced earlier in [Section 2.1](#), may bring about explosive growth in any field, we will limit our focus to intelligent transportation systems (ITS) in this book. In essence, ITS is a mine, which does not yield any material product such as coal, gold, or diamonds and also does not have a bottom. It is a fountain of abstract intellectual ideas that can amaze and transform the world for the better. The recognition of the ITS mine is timely, given that our civilization’s aspiration is rapidly shifting from the material to the abstract. The idea of building tall skyscrapers and long bridges no longer captures our imagination. We have been there, done that, and there is nothing new. We are looking for new challenges that will help us grow; and these lie primarily in the abstract world. One example that probably sounds ridiculous today but will surely come to fruition in the future is the desire to carry out a meaningful conversation with Panini on why and how he invented Sanskrit grammar so that we can design a whole new language for a whole new kind of computer. Another example, equally ridiculous today, is to inquire of Archimedes what were his reasons to invent the odometer and lay down mile markers along Roman highways, beyond the obvious goal of planning troop movements and logistics to quell rebellions in the Roman Empire. A third example is our desire to genuinely discuss with Thomas Jefferson and obtain his insights regarding some of the contemporary troubling issues involving the Bill of Rights. In composing the Bill of Rights, what were his and his colleagues’ concerns, fears, joys, and hopes for the future of United States? Such insights can be invaluable for a brand new nation that may be about to come together in the near future. We have in mind, of course, sophisticated reasoning techniques, new computer memory models, nonlinear left and right brain interactions, and modeling and simulation; not time-travel back into the past.

Our implements for mining will include the familiar principles, concepts, theories, and laws from science; mathematics; computing; electrical engineering; biology and medicine; mechanical, civil, and aeronautical engineering; economics, business, and accounting; and humanities, arts, and philosophy, with the exception that they will be blended and synergized in never-before-seen ways. Our mining techniques will be built on the foundation, namely, that intelligence and creativity are fractal; the deeper we dive, the more subtle possibilities, finer knowledge, and intricate subprinciples appear before us, ad infinitum; the higher we climb, the more encompassing, transcendental, increasingly holistic, and vastly expansive principles arise in our mind,

ad infinitum. We characterize this new approach as the transdisciplinary research and education mindset. While a detailed treatment of the transdisciplinary paradigm is beyond the scope of this book, we present here a few examples of how the approach is likely to work. The reader is also referred to References 6, 7, and 8 for additional relevant discussions.

Consider the following under Example 1. Although J.C. Bose is well known for his pioneering discovery of wireless [9] in 1895, less is known about his contributions that had set the stage for the invention of the transistor at Bell Labs, some 60 years later. Bose had become fascinated by Maxwell's postulate of the existence of high frequency electromagnetic waves and Hertz's demonstration of their creation, in the 10–60 GHz microwave range, but little was known about how to detect their presence. Today, every mobile cell phone has a microwave detector, but back in the 1890s microwaves were a sheer wonder. Bose pioneered an unprecedented approach to detect microwaves at 60 GHz. He fabricated junctions consisting of different metals and materials, wired them into an electrical circuit, and connected it to a very sensitive galvanometer. When he exposed the junctions to microwaves, the circuits registered minute current traces. His sensitive instrument design had been spurred by his scientific curiosity whether plants react similar to human beings and may be considered "alive." While the history of science does not reveal why he chose this approach, a logical explanation is presented here. The use of bimetallic strips had been known from time immemorial and it was widely used for constructing robust chariot wheels. The two metal strips expand and contract differentially when exposed to heat and cold. Bose, a genius, must have argued within himself, as follows. No two materials have the exact same property. Since metals conduct electricity and given that electromagnetic fields and currents relate to one another, it is very likely that the constituent components of a junction will exhibit differential electrical behavior when exposed to microwaves. Clearly, Bose was a transdisciplinary scientist, perhaps one of the first ones in the modern era. He had fused in his mind concepts and principles from electromagnetics, radiation, electricity, chemistry, metallurgy, and plant physiology.

Example 2 focuses on an unusual problem [10,11]. In the mid-1990s in California, State highway 71 needed repair. Traffic would have to be diverted from highway 71 onto another highway, State Route 83, for the duration of the repair, estimated at 6–12 months. The maximum estimated increase in the traffic volume in the afternoon was from 580 to 1460 vehicles per hour. This would increase the noise level by 3 to 4 dB. The U.S. Fish and Wildlife Service voiced their concern that an endangered songbird species, least Bell's vireo, that lives along Route 83 on a stretch of 200 feet would be adversely affected by the increase in the ambient noise level. The adult birds would probably be scared away, leaving the eggs and young chicks vulnerable, eventually jeopardizing the entire population. Also, the increased road noise was thought to interfere with the auditory signals used by the birds to warn each other of danger, detect predators, and monitor their young. The U.S. Fish and Wildlife Service and the California Department of Transportation (DOT) were deadlocked. A brilliant individual, with a deep understanding of signal (S), noise (N), and the true meaning

of the S/N ratio, reflected on the problem and proposed a unique plan. The individual reasoned that, despite the presence of the ambient noise (N), the songbird is able to survive, implying that the S/N ratio, where S represents their auditory signal strength, is well within the limits of tolerance. Clearly, highway construction may proceed if the S/N ratio can be maintained at the present level. The following plan of action was developed. First, the current ambient noise level along Route 83 would be measured and recorded, say, N1. Clearly, the bird species had learned to tolerate N1. A few years earlier, in 1987–1988, biologist John Rieger had computed the tolerance level at 60 dB. Next, the highway department would reduce the posted speed limit from 55 to 35 mph; erect noise barriers along Route 83 to abate the increase in the road noise level; and continuously monitor the noise level on the other side of the barrier, as traffic is diverted away from State highway 71, making certain that it never exceeds N1. The plan worked beautifully.

The third example represents an experiment that had been carried out in a classroom setting in Arizona. First, a group of student were provided a number of reading assignments, one of which related to Example 2. In the final take-home exam, the students were assigned the following problem, a real-world anecdote. The Kansai airport [12], built by Japan, is an artificial island in Osaka bay, deliberately located faraway from the Japanese mainland so that the airport could be kept operational 24 hours, 365 days a year, without driving the inhabitants of Osaka deaf or to insomnia. The decision to locate the airport so far away from the mainland implied that underwater construction would have to confront both alluvial and diluvial soil. While the properties of alluvial soil were well understood, especially when subject to heavy loads, those of the diluvial layer were relatively unknown. As construction proceeded, significant problems surfaced. The entire island was sinking faster and deeper than expected and, more critically, the sinking was uneven across the island. The net result was cost overrun of several billion dollars, a delay of a year in opening the airport, and the precipitous fact that the island is continuing to sink into the ocean, even at the present time. The students were asked to examine every available document on the topic, which in itself was a challenge since the Japanese government had been attempting to systematically erase all available records to avoid public panic; study the problem of sinking; reflect deeply on the fundamental issues; and present recommendations for both corrective actions as well as design guidelines for building future island-airports. The students returned with a simple yet remarkable recommendation. As in the example of the endangered bird species along State Route 83 in California, the Kansai project team should have started by first measuring the ambient noise generated by airports for different traffic volumes and its decline as one moves outward from the airport, in a radial direction. Second, the average noise level that is deemed tolerable by people, during the day and at night, should have been determined through a combination of actual measurements, surveys, and computer modeling. As an example, New York state law categorically limits any noise level to 90 dB. Armed with these two sets of measurements, the project team should then have calculated how far out into the sea should the airport be located. Conceivably, the airport would have been placed lot closer to the mainland and the problem of confronting the diluvial

layer could have been avoided completely. While this recommendation may appear simple, in hindsight, it is not trivial. For, otherwise, the Japanese team would have certainly utilized it and avoided the incredible difficulties they had encountered.

To gain insight into possible ways in which we can exploit the ITS mine, consider the following example, also a true anecdote. If one drives to the Dallas–Forth Worth airport to receive a visitor arriving from abroad, he would normally park in the international arrivals terminal (D) and walk up to the lounge, where arriving passengers would be led after they had successfully completed the customs and immigration formalities. Just outside the set of doors leading into the lounge from the customs and immigration hall is an electronic information board displaying information on the international flights and their scheduled times of arrival. The lounge is full of friends and family members who have come to receive their friends or loved ones. After a little while, many of them become very impatient; sufficient time has passed since the scheduled flight arrival time but there is no sign of their friends or loved ones. The electronic board never cares to display whether a given flight is delayed or landed, information that is readily available to the airport authorities and accessible at the respective airlines' web site. Consider a second example, also a true anecdote. An individual, Mr. Jones, takes a flight from London to New York and lands at JFK airport after seven tiring hours. Mr. Jones is eager to go home in the city but, unbeknownst to him, it is snowing hard. He rushes through customs and immigration, runs outside the terminal to get a taxi, and then spends the next eight hours on the road, while the entire city is paralyzed in a gridlock. Having just arrived at New York, it is not reasonable to expect Mr. Jones to be aware of the driving conditions in the city; however, the information had been known to key transportation officials. If only Mr. Jones was updated on the situation, he would have probably checked in at a hotel within walking distance, spent the night at the airport lounge, or he might have simply sat down for a long dinner at the airport, and certainly refrained from contributing to the gridlock. These two simple examples reveal a serious underlying problem. Today, information technology (IT) is quickly degenerating into a mindless information-providing service. There is little to no thought invested into how the information could be best used by the target audience, let alone if it may mislead the audience or be misused. This, in turn, opens up a tremendous opportunity, namely, develop an invaluable collaboration between transportation officials, who are familiar with the meaning of the traffic-related information, and IT specialists, who are trained to access the most recent data and can determine the best form in which to present the information timely, and to those who need it the most. The final outcome can be highly beneficial and constitute genuine service to society.

## **2.3 EXAMPLES OF FORMIDABLE CHALLENGES AND AMAZING OPPORTUNITIES**

This section presents a few of the ITS problems, under a vastly expanded scope of ITS, that may be justifiably labeled very challenging to nearly impossible. In the true spirit of human innovation, they most certainly can be solved. Furthermore, if and

when they are successfully conquered, unprecedented benefits will flow from them. The intention of this section is to inspire and guide us in the right direction.

- All gasoline and diesel engines, manufactured and in use today, whether in automobiles, diesel locomotives, airplanes, or consumer electric generators, are based on the principle of controlled explosions, which was discovered more than a hundred years ago. Except for new materials including alloys and a limited few improvements, stemming primarily from the use of electronics and computer chips, the technology today is virtually unchanged from that hundred years earlier. The current technology reflects only an incremental improvement. Fundamentally, none of the new principles that have been invented and discovered in the past 100 years have been truly integrated and fused to facilitate radically new engine designs. This implies unlimited opportunities for unprecedented new engine designs of exceptional capabilities, with entirely new attributes and quantum jump in efficiencies.
- From ancient times, bridges have been built to span streams, canals, rivers, and even mini-oceans. While the scientific bridge design principles have not changed over the years, our engineering understanding has steadily improved to the point that today much of the analysis may be carried out by automated computer programs, called finite element analysis. Recent efforts include the use of strain gauges on bridges with appropriate data collection and transmission technology as dynamic weight stations (Prof. Chelsea White, pers. comm., 2010). To date, however, all of the bridges, small or large, wooden or steel, suspension or otherwise, share one common and easily visible attribute. They are all stagnant, stationary, and passive in the sense that they are inanimate, that is, they cannot react dynamically to changing environmental parameters. While it may sound astonishing, there is absolutely no logical reason why we cannot impart intelligence to bridges, which will enable them to sense and adapt to rapidly changing circumstances and protect, up to a point, themselves and their occupants or dependents. As an example, consider that in many of today's computer processor chips, a sensor continuously monitors the residual heat in the core during operation and in the event the temperature exceeds a predetermined threshold the CPU shuts itself down. The incorporation of intelligence helps the computer chip save itself from catastrophic destruction.
- Electric cars, especially hybrid cars that feature both gasoline engines and electric drives, appear to be gaining popularity because of their practicality. A major advantage of these cars is their ability to limit the pollution in highly congested metropolitan cities. During stop and go driving, which is common in urban environments, the use of electric drives eliminate emissions from prolonged periods of idling among conventional gasoline engines. Limited field assessment tests in Colorado (Jim Himelick, pers. comm., 2009) show that plug-in hybrid electric vehicles (PHEV), when charged overnight from the electric power grid, can considerably reduce greenhouse gas

(GHG) emissions [13,14]. Unresolved issues relate to variability in electric consumption and the difficulty in directly connecting fuel and electricity consumption. For non-plug-in hybrid electric cars, however, the ultimate source of energy is still gasoline, while that for battery-operated electric cars is electricity from the power grid, whose primary energy source may be fossil fuel. In essence, these vehicles help redistribute the pollution, away from urban areas and into rural areas where the power-generating plants may be located. The net pollution remains unchanged. Of course, if the world's electrical energy could be derived entirely from nuclear and hydroelectric sources, the emission of carbon dioxide, carbon monoxide, sulfur dioxide, and other poisonous gases from burning fossil fuels would be eliminated. Thus, the challenge to design vehicles that can lower the total emissions continues to remain at the forefront.

- Our lack of precise depth perception is well documented in the field of psychology and is a major factor behind many accidents, especially involving vehicles and locomotives. A large percentage of vehicular accidents that occur at night and that involve the elderly may also be attributed to our inability to precisely gauge the distance of an object directly ahead of us, either receding away or approaching toward us. Along very busy roads in highly populated cities around the world, a primary cause of accidents between vehicles and pedestrians trying to cross the streets, especially at night or when visibility is impaired, may be traced to this flaw in our vision system. As the world's population continues to age, the problem will only get worse, implying an enormous challenge to study and mitigate the problem.
- Roads have been built since the dawn of civilization. The roads in the Roman Empire had extended hundreds of miles, the Autobahn in Germany spans thousands of miles, while the Interstate in the United States runs for tens of thousands of miles. While the building materials and style of construction has evolved over thousands of years, paved roads, by definition, have been built uniformly throughout their entire length. The millions of miles of railroad tracks, worldwide, are also built along the same principle. Naturally, the cost of building the roadway infrastructure is computed and expressed per unit length, and the cost today is staggering. Cost estimates for U.S. Interstates range from a low of \$1 million per mile [15] to \$1 billion per mile [16] for a four-lane highway rated for a 20-year life. Of much greater importance, the cost of maintenance has escalated beyond reach, and even the United States is faltering. As a result, 50 years after being launched, the current U.S. Interstate roadway system is beginning to crumble [4]. In reality, as they age, concrete pavements develop tiny cracks, through which water from melting snow and rain seeps in and amplifies the cracks. This leads to bigger and wider cracks with the passage of time, ultimately disintegrating the pavement into smaller pieces of harmful debris and flying missiles. Lacking the enormous funds necessary to completely overhaul the

thousands of miles of pavements, transportation officials are choosing to superficially repair select sections of the pavements under the “patch and pray” mode. The situation is very serious, for it threatens to undermine the nation’s ability to compete in the world. The challenge confronting us is whether starting from basic physics and mathematics and employing all of the scientific advances we can conceive a much deeper principle and achieve astonishing efficiency in pavement design, namely, reduce both the cost of building and maintain the roadway infrastructure by an order of magnitude or better while greatly lengthening the life. Following in the footsteps of the United States, both India and China are presently engaged in massive construction efforts to expand their roadway infrastructure at enormous expense and will very likely face the same problems 50 years into the future. They would benefit immensely from a radically new thinking at this early stage of their projects.

- While cars, in general, are built to be driven at sustained speeds of 100–120 mph, the average speed of cars on U.S. highways is generally below 60 mph for a number of different reasons. The U.S. DOT uses a rule of thumb, namely, 10,000 cars/h for a four-lane highway, which yields, assuming a three-car length safety rule, an underlying average vehicle speed of 30 mph. On one hand, higher speeds imply shorter travel times, which imply a competitive edge for business and also greatly increased quality of life. On the other hand, posted signs along U.S. highways, “speed kills,” reminds us of a simple physics law that the probability of being seriously injured or even killed increases with speed. In the United States, speeders are routinely prosecuted and fined and, in extreme cases, arrested for driving too fast. The Autobahns in Germany have no posted speed limits, citing the greatly improved pavement conditions, strict driver training, and superior car safety parameters. Clearly, speed and safety issues appear to oppose each other. The challenge for us is to come up with new thinking and extraordinary design principles that will permit cars to travel, no less safely than today, at speeds of 250 mph or more. Much greater distances, say a 1000 mile trip, may be traversed in a reasonable 4 hours, implying exceptional accessibility to the farthest corners of a contiguous land mass and a truly viable alternative to airline travel, except for those routes that require ocean crossing.
- It is common knowledge that improved braking, especially on slippery and icy pavements, may significantly lessen the impact of many accidents involving automobiles, trucks, trains, and even planes. For some of the high-performance racing cars, under normal dry pavement conditions, the powerful brakes can bring the vehicle to a complete stop from a speed of 120 km/h in approximately 60 meters or 10 seconds. A formidable challenge would consist in developing a radically new braking principle, based on an extraordinarily deep understanding of the principles of deceleration, that would reduce the braking distance to a mere 5 meters or stopping time of



1 second, while ensuring the same degree of safety for the occupants as in the present cars. This would imply an order of magnitude improvement and may even preempt many accidents.

- The effects of vehicular accidents throughout the world range from fender benders to minor injuries, long-term medical incapacitation, and, sometimes, painful death. Accidents affect individual and family lives and represent a curse of modern advances. Of far greater importance is the rise of unseen and hidden medical wounds that cannot be reliably detected by the most sophisticated diagnostic medical instruments. While the introduction of safety cages, collapsible engine components, seat belts, airbags, and other advances have done much to increase occupant safety, the number of fatalities, according to a senior U.S. DOT official, continues to hold at 40,000 annually. A key contributing factor is collisions at the street intersections. Furthermore, with respect to pedestrian safety, the commensurate improvement has been slow. The challenge for us is to understand the cause and nature of accidents at a much deeper level and yield revolutionary new techniques to mitigate or preempt them. As starters, we may begin with the idea that accidents represent a violation of a basic physics principle, namely, no two pieces of matter can occupy the same space at the same time, and contemplate on the fact that this law applies only to matter, not energy.
- Given that well over two-thirds of the world is covered by water, boats, aquatic vessels, and ships have been around forever. Their travel speeds have generally ranged from 10–15 knots in the ancient times to 20–40 knots in modern ocean liners and nuclear powered battleships and submarines and upwards of 70–80 knots in specially designed, small, high-speed boats. A challenge that appears impossible would be to conceive an entirely new thinking and bring to fruition ocean going vessels rated at speeds of 400 knots or higher. The benefits would be so immense that there may arise a fundamental shift in our mindset. Access to virtually every corner of the earth, in reasonable time, would become an everyday reality. Given that vast areas of the earth are covered by oceans, most of the population centers located by the water, and that most of the world's waters are interconnected, transportation routes and even our living environments are likely to incur a qualitative change.
- Although the basic idea of personalized rapid transit (PRT) first surfaced in the 1970s to 1990s [17,18,19] and exploratory studies had been carried out in Sweden and the United States [20], PRT has not yet been embraced by policy makers and transportation officials around the world. Under PRT, upon arrival at a city, a traveler climbs onto a two- or three-seater private automated car at designated stations at the airport and then punches in his or her desired destination from a series of choices around the city including hotels, convention center, shopping malls and eateries, museums, government complexes, and major downtown office buildings. The car



glides on an above-ground guideway and automatically and safely switches within the loop guideway to reach the destination station. Following the traveler's disembarkation, the car may either be immediately reutilized by another traveler already waiting at the station or the empty car may be led to the nearest depot under algorithmic control. PRT promises to greatly reduce congestion in highly populated urban centers. The challenge has been the lack of practical PRT architectures and, more importantly, the absence of a scientific approach to assist policy makers in systematically analyzing the needs of a PRT network, determining whether or not to build PRT networks in a given environment, and, where affirmative, to lay down detailed engineering parameters and synthesize the underlying asynchronous decentralized algorithms for coordination and control.

- During their initial deployment in commercial aircrafts, black boxes had been invaluable in determining the cause of incidents and accidents. This helped greatly in incorporating appropriate safety features and redesigning the aircraft controls and protocols. Lately however, black boxes from airline disasters seem to be less and less useful. The problems may be organized into two broad categories. First, where airline disasters have occurred over the world's oceans, the black boxes tend to sink deep into the abyss, defying easy retrieval. Given that two-thirds of the earth is covered by water and that many of the international routes involve crossing the oceans, new thinking is needed in black box redesign. Second and more important, many of today's disasters appear to occur so abruptly and quickly that either the black box data is too little, the information is less than valuable, or the cockpit voice recording data is utterly irrelevant to the accident investigation. It appears that the crew had very little time to react to the threat. Unlike in the past, today's aircraft controls are significantly more sophisticated, stemming from the use of fast computers and complex software. When a problem causes a plane to crash, despite the fast computer controls, one thing is clear. The threat is extremely complex and is beyond the realm of scenarios that logically flow from the present understanding of the technology. There is also a strong possibility that the current black box designs are unable to identify and record the right set of aircraft operational parameters at the correct resolution. The challenge for us is to develop a whole new approach to aircraft safety, focusing more on rapidly detecting and automatically neutralizing the threat while deemphasizing the tradition of capturing flight data and cockpit voice recording for after-the-fact analysis.
- An excess of water at the wrong place and wrong time can cause significant damage to life and property. When a river continues to rise because of excess rainfall upstream, it may begin to threaten the levees and embankments. When water accumulates from flooding, excess rainfall, storm surges, or failed levees, it can destroy homes, crops, communities, and our very livelihood. Under these circumstances, we have relied on the mercy of

gravity and other natural forces for the past thousands of years. A formidable challenge for us is to conceive an entirely new thinking and synergize scientific principles toward an extraordinary approach to drain water away at incredible speeds. The benefits would be indescribable, especially against the backdrop of the irrecoverable damage from Hurricane Katrina back in 2005.



---

# 3 Fundamental Issues in Transportation Systems

By their very nature, transportation systems span nontrivial physical distances. Goods and people are transported from one physical location to another utilizing routes that pass through one or more interchange points, also known as hubs or route points in the literature. The organization of the interchange points facilitates the sharing of resources and offers other conveniences such as refueling, and so on. Thus, the constituent entities of transportation systems include the goods and people being transported and the interchange points, all of which are geographically dispersed. Prior to the discovery of electromagnetic communication, information on the transport of goods and people was propagated along with the transported material itself or was physically carried by some other means, the speed of which was similar in scale to the rate of movement of the goods or people. Thus, in essence, the transit speeds of the good and people and the information about their transport were similar.

The first major revolution in transportation coincided with the introduction of electromagnetic communication, which enabled the propagation of information about the movement of goods and people significantly faster than the actual transport of the material at limited speeds. Thus, a modern transportation network may be viewed conceptually as consisting of two principal components: an information network that transports pure electromagnetic energy and a material transport network that carries goods and people. While the information network is optional but highly beneficial, the material transport network is a fundamental requirement.

The availability of computing engines fueled the second major revolution in transportation systems wherein fast and precise computers were exploited to efficiently control and coordinate the transport of goods and people across the system. Excellent examples of computer usage in transportation in use today include the centralized control for railways [21], centralized air traffic controllers, and traffic management centers [22] for automobiles. For simplicity and ease of comprehension, the coordination and control functions, from the start, have been consolidated into centralized units, which continues to dominate to this day. The significant cost and bulk associated with the earlier computers had reinforced the decision to centralize. Under the centralized paradigm, first relevant information is acquired from every constituent entity of the system at the centralized unit. The information may include the origin and target destinations of the units being transported; the transport-related preferences, if any, of every unit; the state of occupancy of the transport paths between the interchange points; and the state of the information network. Next, the centralized unit executes a complex, decision-making program that aims to achieve overall efficiency, subject to specific preestablished criteria. The program computes the movement-related decisions for each of the units that needs to be transported, one

at a time. Then, the decisions are propagated to the respective units, which, in turn, realize them.

Conceptually, the generation of the decisions by the centralized computing engine may be slow, especially where the number of units in the system is large. Furthermore, the tasks of reading the information from all of the remote units and relaying the respective decision to each of the distant units, one at a time, may greatly diminish the speed of decision making as well as the realization of these decisions. Slower decision making implies inaccuracy and imprecision. Furthermore, centralized units are highly susceptible to natural and artificial disasters. In reality, many of today's transportation networks are witnessing a dramatic increase in the number of the constituent units, faster rates of movement of goods and people, and a sharp increase in the demand for faster and more precise coordination and control. As a result, it has been projected that for some transportation systems even the use of today's fastest supercomputers may not be adequate. Recently, the U.S. Federal Aviation Authority (FAA) has initiated the concept of "free flight" [23] to enhance the safety and efficiency of the National Airspace System (NAS). The concept moves NAS from a centralized command and control system between pilots and air traffic controllers to a distributed system that allows pilots, whenever practical, to choose their own routes, dynamically, for efficiency and economy. To achieve safety within "free flight," the FAA must establish guidelines to control the flight routes. Fundamentally, the requirement in the current centralized approach that the decision making for the constituent units, X, Y, . . . , be sequential, is unnecessary. Where the resources pertaining to the transit of X are relatively independent of those pertaining to Y, the decision making for X and Y, in theory, may occur independently and simultaneously.

Today, the transportation discipline is on the brink of experiencing the third major revolution and possibly the most complex—the transformation from the centralized paradigm to the asynchronous, distributed paradigm that will integrate fast computers and high-performance computer networks through novel computer algorithms. The discipline of intelligent transportation systems (ITS) [24] encompasses all of the advances in transportation and calls for the design of innovative and creative approaches to the transportation needs, utilizing the fundamental principles, followed by thorough scientific validation.

### 3.1 PRINCIPAL CHARACTERISTICS OF ITS

The principal characteristics of ITS, which promises to gain increasing importance in the future, include the following:

1. **Information dispels and wears out ignorance:** The single biggest cause of driver impatience that leads to reckless driving, incidents, road rage, and accidents is ignorance. When a traveler encounters an obstacle and does not know what lies ahead, a rational or irrational fear may arise in that he or she might not succeed in reaching the destination on time, which, in turn, may have untold consequences. The ITS architecture's greatest contribution

may consist in providing logically relevant and accurate information in a timely manner to needy travelers in a form that they can utilize to determine alternatives and recompute their plans. For example, consider that a busy executive, say Joe, in some city, scheduled to be on a flight, say XYZ, from a nearby airport, is late in leaving the office. Joe panics and drives reckless, nearly running over a pedestrian, and barely makes it to the airport in time to catch the flight, only to find out that XYZ has been delayed for 2 hours since the incoming aircraft had developed an unexpected mechanical problem. If under ITS, the flight delay information, as soon as it becomes available, is automatically transmitted to every passenger scheduled for XYZ, in this case Joe's cell phone, most if not all of Joe's driving-related problems may disappear immediately.

2. **Automated computation:** Unlike in many transportation systems across the world where the decision making and the computation of the arrival times are still estimated manually, future system architectures must employ automated decision-making computer systems to yield accurate information and achieve precise control and coordination.
3. **Accuracy of information is crucial:** Information must be accurate, timely, relevant, and consistent. Otherwise, drivers will begin to question the value of the ITS technology and may even abandon it prematurely. Any out-of-date piece of data can quickly assume the form of misinformation with severe unintended consequences. Consider a recent event, intended to serve as an example of what must not occur under ITS. As recently as in June 2002, a programmable overhead electronic information board on Interstate-10 warns drivers that the air quality between Phoenix and Tucson is poor. The driver reading the sign finds it informative but irrelevant since he or she is driving west toward Los Angeles and is just about to leave behind the Phoenix city line. The information should have been displayed on the opposite side of Interstate-10 between Phoenix and Tucson.
4. **Respect for drivers' freedom is of paramount importance:** Except under true emergency conditions, the ITS architecture should neither attempt to control nor dictate a driver's behavior for two reasons. First, no centralized authority [25] can ever know with certainty the goals, objectives, and thinking of every driver and, therefore, any attempt to control him or her will invariably be based on erroneous assumptions. Second, in the long run, drivers will resent the intrusion on their freedom to make their own decisions and will ultimately abandon ITS. The key principle is to provide drivers with as much relevant information as possible and equip them with appropriate networking and computing resources so that they can determine and freely choose the right course of action under the given circumstances.
5. **The demand for flexibility and freedom of choice:** The frequent lack of flexibility, the absence of personalized services, and the availability of mostly inaccurate estimates are increasingly being rejected by the customers of transportation systems. There is strong unwillingness to accept the

traditional excuses of limited computational ability and network bandwidth. There is an increasing demand for flexibility, freedom of choice, and personalized service, and the trend is likely to continue into the future.

6. **The demand for accurate, that is, precise and up-to-date, information:** In the traditional approach, data are first collected at a centralized unit, processed, and the resulting information is disseminated to the geographically dispersed customers. Given the geographical distance and the finite speed of propagation of electromagnetic radiation, when a customer intercepts information relative to the transport of a unit in transit, a finite time interval has elapsed since the information was originally generated. For dynamic systems, this delay implies that the information received by the customer has incurred latency and is, in essence, inaccurate and imprecise. The degree of the error due to latency is a function of the length of the delay, relative to the dynamic nature of the system, and the resolution of accuracy. Thus, latency is fundamental to every transportation system, and future system architectures must focus on distributed schemes that aim at eliminating all unnecessary sources of latency, where possible, and realizing efficient, accurate, and timely decisions.
7. **The individual traveler must constitute the focus:** The ITS architecture must be fundamentally centered around each individual driver or traveler, subject to safety and fair resource availability for all. That is, while gross metrics such as pavement utilization, average travel time, network bandwidth usage, and compute horsepower are important to the planners and operational managers, concern for the individual traveler's parameters must remain at the forefront.
8. **A fundamental, enabling characteristic of transportation networks:** Although it is an obvious fact that matter in the form of goods and people is transported in a transportation network, the implication is profound. The material units being transported may carry with them their own computing engines, which, at the present time, are necessarily matter. In contrast, in a communications network, the units of information constitute pure electromagnetic energy, and they cannot carry with them their own computing engines while in transit. While carried along with the goods and people, the computers, in turn, may facilitate dynamic, travel-related computations and decision making. Also, since the constituent units of a transportation system may communicate with each other while in motion, utilizing wireless or infrared techniques, there is hope that the need for centralizing the information gathering and decision-making functions may be eliminated. The trend of decreasing physical size and cost, increasing capability, and lower power consumption in computer designs is encouraging and will likely render their use in transportation networks increasingly practical.
9. **The design of asynchronous, distributed algorithms for control, coordination, and resource management:** Since the constituent units and the resources of any transportation system are geographically dispersed, it is

logical for future system architectures to exploit distributed algorithms. The units to be transported across a network are likely to request for service, independent of one another and at irregular intervals of time. Thus, the interactions in the system, in essence, will be asynchronous, requiring the design of asynchronous, distributed algorithms for control, coordination, and resource allocation.

By design, an asynchronous, distributed, algorithm for a transportation system must necessarily reflect the highest, meta-level purpose or intent of the system. The algorithm manifests itself in the behavior of every constituent unit. It will hold the potential of exploiting the maximal parallelism inherent in the system. Furthermore, local computations must be maximized while minimizing the communications between the entities, thereby implying high throughput, robustness, and scalability. The key properties of such algorithms are as follows:

- a. **Identification and definition of entities:** From the perspective of control, coordination, and resource allocation algorithm, entities constitute the basic decision-making units and define the resolution of the decision behavior of the transportation system. From the perspective of the physical system, entities correspond to its natural, constituent elements and include the resources and the units to be transported. An entity must be self-contained, that is, its behavior, under every possible scenario, is completely defined within itself. Every entity exists independent of all other entities and, therefore, its behavior is known only to itself. Unless the entity shares its behavior with a different entity, no one has knowledge of its unique behavior. Conceivably, an entity will interact with other entities. Under such conditions, its behavior must include the scope and nature of the interactions between itself and other entities.
- b. **Asynchronous nature of entities:** In general, units requesting transport are incident on the system at irregular intervals of time. A unit may consist of a single traveler or a group of individuals traveling together under a special event. Since the units are incident at different geographical points of the system, they are unaware of the presence of each other, their rates of progress are likely to be different, and their destinations are likely uncorrelated and different, the interactions in the system will be asynchronous. Time constitutes an important component in the behavior of the entities and their interactions. Although every entity, by virtue of its independent nature, may possess its own unique notion of time, when a number of entities  $E_1, E_2, \dots$  choose to interact with each other, they must necessarily share a common notion of time, "universal time," that enables meaningful interaction. The universal time is derived from the lowest common denominator of the different notions of time and reflects the finest resolution of time among all of the interacting entities. However, the asynchronicity manifests as follows. Where entities A and B interact, between their successive interactions, each of A and B proceeds



independently and asynchronously. That is, for A, the rate of progress is irregular and uncoordinated (with respect to B) and reflects lack of precise knowledge of the rate of progress of B, and vice-versa. At the points of synchronization, however, the time values of A and B must be identical.

- c. **Concurrency in the entities:** Since every entity exists independent of others, except for the necessary interaction with other entities, every entity must necessarily be concurrent. That is, the progress and rate of its execution is independent of those of other entities and, at any given instant during execution, the states of the entities are unique. In an implementation of the algorithm, every entity may be mapped to a concurrent process of a host computer system. The recognition that entities are concurrent is important for two reasons. First, it reflects the correct view of an actual transportation system and, as a result, the control and coordination algorithm represents reality closely. Second, when the algorithm is executed on a host computer system that contains adequate computing resources, the concurrent entities may be executed simultaneously by dedicated computing engines to achieve faster overall execution. Faster execution will enable, realistically, the execution of the algorithm a large number of times, for different parameters, yielding insight into the system design issues.
- d. **Communication between entities:** Clearly, an entity may not possess total and accurate knowledge of everything it may need for its continued functioning at every instant of time due to the geographical distances and the finite speed of propagation of information. Therefore, the sharing of data and knowledge through interaction with other entities may constitute an important and integral component of the algorithm. Entities may interact with one another, and the nature of the interaction may assume different forms. First, each set of entities that interact between themselves are identified, and they reflect the corresponding real-world system exactly. In an extreme case, any entity may interact with any other entity. Second, the necessary data and information must be shared between the interacting entities and appropriate message structures must be developed. Third, the information may be shared on a need-to-know basis, to ensure privacy and precision. Fourth, all message communication is assumed to be guaranteed. That is, once a sender propagates a message, it is guaranteed to be received by the receiver(s). Thus, the sharing of data and information among entities implies a communication network that interconnects the entities and is an integral component of the algorithm. The topology of the network is determined by the nature of the interactions.
- e. **Proof of correctness of the algorithm:** By their very nature, the control and coordination algorithms, involving hundreds of autonomous entities executing simultaneously and asynchronously with no centralized agent responsible for controlling their activities, are difficult for the sequential

human mind to comprehend. To ensure accuracy, correctness, and safety of a real-world transportation system under algorithm control, it is therefore crucial to develop a proof of correctness. Fundamentally, the proof of correctness must guarantee that the system operates correctly, that is, the execution of the algorithm must accurately reflect reality, there is the absence of any inconsistency such as accidents, and that the system progresses toward its unique objective. The proof of correctness is especially important since each decision-making entity utilizes a subset though a relevant fraction of the system-wide information and the data from other entities is subject to latency.

- f. **Robustness:** The asynchronous, distributed algorithm is expected to yield a robust system, which, unlike a centralized system, is much less susceptible to natural and artificial disasters. Each geographically dispersed entity is a decision-making unit, and the likelihood of a disaster affecting every entity is remote. Thus, even if one or more of the asynchronous entities fail, the remaining entities are not likely to be completely affected and the system is likely to continue to function in a degraded mode. For the algorithm to operate under partial failures, exceptional handling must be incorporated into the design of the entities and their asynchronous interactions.
- g. **Performance:** The use of multiple processors executing concurrently under algorithmic control implies superior performance. The degree to which the algorithm is able to exploit the parallelism inherent in the underlying system is reflected in its performance metric. While every specific transportation system is likely to require its unique set of metrics, two criteria may be applied uniformly across all algorithms. First, “performance scalability,” detailed in Chapter 4, reflects the ability of the algorithm to continue to function, that is, achieve the primary performance objective, despite increasing system size, that is, increase in the number of entities constituting the system. Since the computation underlying the system-wide decision making is distributed among every entity, as the system size increases, both the demand for increased computational power and the number of computational engines increases. Assuming that the communication network experiences proportional expansion, the ratio of available computational power to the required computational power is expected to decrease only by a marginal fraction, implying that the achievement of the primary performance objective will be affected only marginally. Second, consider a hypothetical mechanism that is capable of determining the absolute performance of any given real-world problem, which, in turn, may serve as the ideal metric against which the performance of any algorithm may be evaluated. Chapters 5 and 6 present this metric in detail. Performance measures can play an immense role in investment and operational decisions of transportation agencies.

- h. **Stability:** Where the algorithm constitutes an implementation of a real-world transportation system, it is likely to be subject to unexpected changes in the operating conditions. The property of stability refers to the behavior of the algorithm under representative perturbations to the operating environment and is detailed in Chapter 7.
10. **Simulation is indispensable:** For many of today's transportation systems, given the increasing size and complexity, which implies a large number of variables and parameters that characterize a system; the wide variation in their values; and the great diversity in the behaviors, the results of the analytical efforts have been restrictive. ITS systems are likely to be complex, implying that modeling and large-scale asynchronous distributed simulation may be the most logical and, often, the only mechanism to study them objectively. Key benefits of modeling and simulation are many. First, they enable one to detect design errors, prior to developing a prototype, in a cost-effective manner. Second, simulation of system operations may identify potential problems, including rare and otherwise elusive ones, during operations. Third, analysis of simulation results may yield performance estimates of the target system architecture and potential for growth.
11. **Continual checking for operational errors:** By their very nature, complex systems may contain design errors that manifest irregularly during operations and elude detection but are severely damaging. To address this weakness, ITS architectures must incorporate automatic mechanisms to continually accept any errors reported by users and bubble them up to the system architect level so that corrections and refinement may be realized in a timely manner.
12. **Continual checking for inconsistencies between ITS principles and societal norms:** Any ITS architecture must be based on sound logical principles, and where established norms pose inconsistencies, they must be carefully analyzed. For example, virtually everywhere in United States, high occupancy vehicle (HOV) lanes, where present, are located at the extreme left. Since pavements to the left are designated fast lanes, the immediate connotation is that the HOV is a fast lane. This is reinforced by the fact that, often, fast non-HOV cars use the HOV lane to pass slower cars in the left lane. Furthermore, to get to this lane, an HOV entering the highway from a ramp typically located on the right must maneuver through fast-moving cars in the left lanes, if at all possible on a congested day, leading one to infer an underlying assumption that HOV drivers intend to drive fast. The assumption may be seriously wrong since many of the HOV vehicles may represent families traveling together who would prefer to drive at the speed limit and avoid accidents. As a second example, consider the nature of many of the highways whose origins date back to the 1700s before the union was formed. To drive from town S to T, one is advised to follow highway H, which will lead straight to a major interstate freeway. Although H appears as a straight line road on the map with no major intersecting highways except the interstate

freeway, a driver unexpectedly encounters forks at three different locations. The peculiarity of these forks is that if one follows the most obvious branch of the fork that has the same make and feel of the road that he or she was on that would be a mistake. To continue on H, one would have to make a sharp turn, left or right, on the less obvious branch of the fork. There are signs posted at the forks but worn out and obscured by years of overgrown vegetation. Presumably, pieces of the road date back hundreds of years and when consolidated into highway H the sharp jogs were left in place to avoid running through people's homes, places of worship, and so on. Clearly, thousands of drivers following H must have gotten delayed or lost, and the cumulative delay over the past fifty decades has run into hundreds of thousands of hours. To justify such waste today is very difficult, and for ITS to make a genuine difference, these problems must be addressed scientifically and objectively.

### **3.2 SCIENTIFIC VALIDATION OF ITS DESIGNS THROUGH MODELING AND SIMULATION**

The traditional approach to understanding the behavior of real-world transportation systems has been to develop analytical models that attempt to capture the system behavior through exact equations and then solve them using mathematical techniques. This has been adequate in the past and may continue to serve effectively in many disciplines. However, for many of today's transportation systems, given the increasing size and complexity that implies a large number of variables and parameters that characterize a system, the wide variation in their values, and the great diversity in the behaviors, the results of the analytical efforts have been restrictive. For example, Partners for Advanced Transit and Highways (PATH) researchers, at the University of California, Berkeley, have proposed an architecture for intelligent vehicle highway systems (IVHS) [26], wherein one or more automobiles are organized into discrete platoons that move through special lanes, similar to HOV lanes, on existing freeways at very high speeds. When a vehicle enters into the network and announces its ultimate destination, the system assigns it a nominal route through the network. In their attempt to develop an analytic model of the architecture, utilizing the principles of control theory, PATH researchers encounter a system with a formidable number of states. Tomorrow's systems are expected to be far more complex, implying that modeling and large-scale simulation may be the most logical and, often, the only mechanism to study them objectively.

Modeling refers to the representation of a system in a computer executable form. The fundamental goal is to represent in a host computer a replica of the target transportation systems architecture including all of its constituent components, as accurately and faithfully as possible. Simulation refers to the execution of the model of the target system design on the host computer, under given input stimuli, and the collection and analysis of the simulation results. The benefits of modeling and simulation are many. First, they enable one to detect design errors, prior to developing a prototype, in

a cost-effective manner. Second, simulation of system operations may identify potential problems, including rare and otherwise elusive ones, during operations. Third, analysis of simulation results may yield performance estimates of the target system architecture. Unlike in the past, the increased speed and precision of today's computers promises the development of high-fidelity models of transportation systems, ones that yield reasonably accurate results, quickly. This, in turn, would permit system architects to study the performance impact of a wide variation of the key parameters, quickly and, in a few cases, even in real time or faster than real time. Thus, a qualitative improvement in system design may be achieved. In many cases, unexpected variations in external stress may be simulated quickly to yield appropriate system parameter values, which are then adopted into the system to enable it to successfully counteract the external stress. In simple terms, simulation permits the playing of "what if" games. This will allow designers and operators to conceive hypothetical, yet realistic, scenarios that may arise in the future; develop solutions and strategies; and test them off-line, i.e., non-operational mode, in a cost-effective and safe environment. Last, the design of new performance metrics may be facilitated to gain a better understanding of the nature of the system behavior. The issue of determining appropriate input traffic demand patterns is discussed in each of the Chapters 4 through 8.

---

# 4 DARYN: A Distributed Decision-Making Algorithm for Railway Networks

## 4.1 INTRODUCTION

A railway network typically consists of thousands of miles of tracks and hundreds to thousands of locomotives carrying goods and people from one place to another. For efficiency and modularity, usually, the tracks are divided into individual units, each of which may be controlled exclusively by the system. Thus, a train in propagating from a location A to another location B may travel over several tracks. Given that two or more trains may need to use, at some time instant, the same track and that only one train may occupy a track at any time, the principal goal of the railway system is to allocate tracks to trains such that (1) collisions are avoided and (2) the resources are utilized optimally.

Traditional approaches utilized the principles of centralized scheduling to achieve these objectives. In centralized scheduling, the destination of every train is known a priori by the dispatcher—a uniprocessor computer. In addition, a dispatcher receives, at regular intervals of time, the current position and speed of every train and the status, that is, whether occupied or empty, of every track in the system. The dispatcher analyzes these data sequentially and based on a cost function and computes the subsequent subroute that every train must execute. Thus, the dispatcher controls the status of every track and the subsequent movement of every locomotive to achieve overall optimal efficiency of resource usage.

Consider a railway network, shown in [Figure 4.1](#), where the four tracks  $T_1$  through  $T_4$  connect the four stations A through D. The centralized scheduler is represented through the block DPU in [Figure 4.1](#) that is connected to each of the tracks  $T_1$  through  $T_4$  through solid lines implying permanent communication link fixtures. Each line represents bidirectional flow of information. The track statuses are fed to the DPU, and the DPU, in turn, propagates new instructions to the tracks. Where a train  $R_1$  is propagating from the originating to the destination station, a temporary communication link (perhaps a radio-link) is established with the DPU. This link is canceled following the arrival of the train at its destination. Assume that at some time instant, trains  $R_1$  through  $R_3$  occupy the tracks  $T_3$ ,  $T_1$ , and  $T_2$  respectively. Thus, temporary links, shown through dotted lines in [Figure 4.1](#), represent the bidirectional flow of information between the trains and the DPU. At certain time intervals, each of  $R_1$  through  $R_3$  and  $T_1$  through  $T_4$  propagate relevant information to the DPU.




FIGURE 4.1 Centralized scheduling in traditional approaches.

The time interval is determined based on the length of the tracks, speed of trains and information processing by the DPU, communication link delays, and the desired accuracy of the decisions. The DPU reads the data sequentially, analyzes them, and computes the subsequent tracks for  $R_1$  through  $R_3$  based on a cost function. The cost function may take into consideration the actual cost of using a track, the priority of the train, and future implications. The DPU propagates these decisions to the trains  $R_1$  through  $R_3$  sequentially and permits them to make appropriate movements. This process continues until all trains reach their ultimate destinations when the system is said to have terminated. For simplicity, the choice of the time intervals may correspond to the time instants when all trains have arrived at a station.

A description of the centralized algorithm, using pseudocode, is presented in Figure 4.2.

Graff and Shenkin [27] describe a simulation of a traditional algorithm, subject to a set of predetermined objectives, where the decision subroutine of the dispatcher is executed whenever a train arrives at a node (i.e., an intermediate station). This algorithm is also utilized by the Southern Pacific Transportation Company. The authors note that the CPU time required on an IBM PC remains manageable, provided that the number of tracks and trains is limited to 15 and 25 respectively. Fukumori [28] provides an artificial intelligence approach to address the scheduling problem for a partially ordered set of events utilizing the timetable as a given specification. Recent improvements to the traditional algorithm include sophisticated strategies to communicate train and track information to the dispatcher-computer. Rao and Venkatachalam [29] present an experimental system that uses microprocessors to interlock control points, monitor track circuits, and interact between the track circuit and the locomotive currently occupying the track. The overall objective is to reduce the probability

```

main ()
{
  open file to read information on all stations;
  open file to read information on all tracks;
  open file to read information on all randomly generated train information;

  interval = 1;
  while not (all trains arrive at destinations) {
    for this interval, read speed and position of every train;
    for this interval, read status of every track;
    determine subsequent tracks for every train through cost function evaluation;
    permit the trains to move to the the next tracks;
    interval <- next interval;
    verify whether all trains have reached destinations;
  }
}

```

**FIGURE 4.2** Algorithm for centralized scheduling.

of accidents. Coll, Sheikh, Ayers, and Bailey [21] present a detailed description of the North American advanced train control system (ATCS), where the allocation of tracks to trains is controlled exclusively by a dispatcher. ATCS is novel in its use of transponder-computers that are embedded along the tracks to monitor the necessary data, onboard computers that enforce the allocations issued by the dispatcher, and a sophisticated computer-to-computer communications architecture. Kashyap [30] describes the role of telecommunications in the management of tracks in the context of Indian Railways. Vernazza and Zunino [31] present a distributed intelligence methodology for railway traffic control, which is described as follows: the entire system of tracks is divided among a few decision centers (DCs) exclusively. When a train  $R_i$  occupies a track that is under the jurisdiction of a decision center  $DC_j$ ,  $DC_j$  assumes the role of the dispatcher for  $R_i$ . When  $R_i$  needs to use a track that is controlled by another decision center  $DC_k$ ,  $DC_j$  must negotiate with  $DC_k$  for the temporary privilege of using the track in question. A principal limitation of the approaches in [21] through [30] is that they are limited to centralized scheduling, and, as a result, such systems are bound to be slow for large systems. Furthermore, this may imply that the control mechanism is ineffective as a real-time system. Evidently, the need to send information from every train and track to the dispatcher will imply slow communication. This results from the large signal propagation delays over the links and the time necessary for the dispatcher to read every piece of data. Conceivably, such delays may force trains to unnecessarily idle [2] at stations awaiting dispatcher's decisions or lead to collisions. The approach [30], while representing an improvement over the earlier approaches, is limited in that (1) it resorts to a limited number of DCs and (2) lacks realistic modeling and detailed simulation study.

The remainder of this chapter presents a novel approach, DARYN, that addresses a number of limitations of the earlier approaches. [Section 4.2](#) presents a detailed description of the DARYN algorithm and the issue of modeling on a loosely coupled parallel processor. [Section 4.3](#) describes the implementation of DARYN on ARMSTRONG.



Section 4.4 presents measurements from the modeling and simulation and details an analysis of performance of DARYN.

## 4.2 THE DARYN APPROACH

### 4.2.1 ALGORITHM


An analysis of the centralized algorithm reveals that the primary reason for the dispatcher to compute the routing for every train lies in the fact that the same track may be required by two or more trains at an instant of time. If the decisions for the trains are computed by agents that are completely independent and noncoordinating, this may lead to collisions that are unacceptable. DARYN proposes a solution that attempts to eliminate collisions, or equivalently, maintain consistency in all of the decisions while distributing the overall decision process among concurrently executable, natural entities that execute with a minimum of synchronizations. Thus, under these circumstances, the maximal inherent parallelism may be utilized. The word “natural” refers to the physical entities that are inherent in the system.

In DARYN, the decision process is distributed among the processors embedded in the stations and locomotives of the system. Every train and station executes its share of the decision making concurrently and independently, subject to necessary synchronizations, such that consistency is maintained and the overall system executes much faster. A processor, associated with a locomotive, possesses complete knowledge of the track layout of the entire network. Since the frequency of changes in the actual layout of the tracks is likely to be extremely low, it is realistic for the train-computer to store this information as a part of its initialization. This chapter assumes that a train-computer has accurate knowledge of the entire track layout at all times. Although the electronic storage requirement may seem significant, it is easily and economically achieved through the use of inexpensive USB thumb-drives whose capacities are rated at 32 GB (and increasing). While most of the track layout may be assumed stable, the relatively few dynamic changes in the track layout due to faults, maintenance repair, and others, may be easily read by the trains from the neighboring station-computers. The train-computer,  $TC_i$  for train  $R_i$  is also responsible for evaluating the cost function and thereby determine an optimal route for  $R_i$ . In order to achieve its objectives, the train-computers must interact with the station-computers that control the tracks exclusively. Each track is controlled exclusively by a single station-computer. For obvious reasons, the tracks in the vicinity of a station are controlled by that station-computer.

Initially, every track is unoccupied and every train, located at the originating station, is uncommitted with respect to its route. Following initialization, every train is fully aware of the entire network, that is, the locations and owners of the tracks and every station-computer is aware of the tracks it owns. A train  $R_i$  utilizes its ability to evaluate the cost function and determine an optimal route. Thereafter, it propagates a request to the station-computer that controls the first track in its optimal route. This request is implemented in the form of a message and is sent by  $R_i$  to the current station  $S_j$ , where  $S_j$  is the owner of the track in question, it examines other similar

requests and arrives at a resolution, utilizing the notions of first-come-first-serve and priority. That is, a station determines the lowest time of all times of requests by trains corresponding to a particular track. The request corresponding to the lowest time is approved and all others are denied. The response to the original request is propagated by  $S_j$  to  $R_i$ , where  $S_j$  is not the owner of the track in question; the request is routed to the appropriate station-computer,  $S_k$ , and a response from it is sent back to  $R_i$  via  $S_j$ . Where the train  $R_i$ 's request is approved, the train may physically proceed on the track toward the subsequent station that constitutes a subset of its optimal route. However, if the request is denied, the train-computer must reevaluate the next best route and repeat the process for the first track of this new route; where a train's requests for tracks are repeatedly denied, either at the originating or any intermediate station, the train-computer must continue to issue requests. It is likely that such a scenario refers to a potential problem. In addition, until a train is given permission to occupy a subsequent track, it may not move from the station where it is currently located. This guarantees absolute consistency of decisions. When the train reaches the subsequent station, it repeats the same negotiation process. Every station-computer performs two functions—(1) issuing permission or denials to requests for tracks that it owns and (2) routing requests from trains to appropriate station-computers as well as responses in the opposite direction. This process continues until all trains reach their destinations.

Thus, in DARYN, a train-computer may issue a request only when it is physically located at a station. Conceivably, the train-computer interacts with the station-computer through optical, infrared, or other means, and the communication process is initiated as soon as the train enters the station. Given that trains may travel at a maximum speed of 200 km/h and that the length of a typical station is 200 meters, the minimum contact time between a station and a train may be approximately 3.6 seconds. This is, at least, several times larger than the CPU time (milliseconds) necessary for negotiations by the digital computers underlying the trains and stations and, as a result, in most cases, a train may not even need to physically stop at a station during the negotiation process. Also, a train-computer necessarily recomputes a new optimal path at a station once its request for access of a track, based on an earlier optimal route computation, is denied. Logically, one may argue in the favor of permitting a train to request reservation on two or more consecutive tracks of its optimal route prior to permitting it to move. While this might facilitate efficient long-term planning, it has the potential to degrade overall performance by reserving tracks too soon, thereby locking out tracks and leading to poor routing for other trains. DARYN is restricted to looking ahead and requesting reservation only up to the immediate track. Furthermore although the nature of the cost function may be complex, in this chapter, a simple cost function is assumed. Every track has the same cost associated with it and every train has the same priority. Assuming an X-Y coordinate system for the train network, where a train has to move  $M$  and  $N$  units in the X (horizontal) and Y (vertical) directions respectively, the cost function forces the train to first move in that direction given by the larger of the values  $M$  and  $N$ . When  $M = N$ , either of the X or Y directions may be chosen.



**FIGURE 4.3** The DARYN approach.

A necessary condition in the DARYN approach is that every station and train must possess clocks that conform to an absolute standard clock. At initialization, all such clocks are synchronized and reset to 0. This ensures consistency of timing of all requests, responses, and decisions.

As an example of the DARYN approach, consider the network in Figure 4.3. In Figure 4.3, the stations A through D are connected by the tracks  $T_1$  through  $T_4$ . Assume that the station-computer for A controls the tracks  $T_1$  and  $T_4$  while each of the station-computers for B and D control the tracks  $T_2$  and  $T_3$  respectively. Thus, the station-computer for C owns no tracks. The station-computers communicate with each other through communication links that are laid alongside the tracks. Assume that, at a time instant, the system contains four trains  $R_1, R_2, R_3,$  and  $R_4$ . Trains  $R_1$  and  $R_2$  are asserted at station A at times given by  $t = 0$  unit and  $t = 5$  units respectively. Both are destined for C. Trains  $R_3$  and  $R_4$  are asserted at station C at times given by  $t = 0$  unit and  $t = 5$  units respectively. Both are destined for A. In DARYN, assume that trains  $R_1, R_3, R_2,$  and  $R_4$  determine their optimal routes as  $\{T_4, T_3\}, \{T_2, T_1\}, \{T_1, T_2\},$  and  $\{T_3, T_4\}$  respectively. Thus,  $R_1$  requests station-computer A for reserving track  $T_4$  while  $R_3$  sends a similar message to station-computer B via station-computer D. Both the requests will be granted and the trains will move on their respective tracks. Assume that the length of the tracks and the trains' speeds are such that  $R_1$  reaches station C at  $t = 7$  units and  $R_3$  reaches B at  $t = 4$  units. At this point, that is,  $t = 4$ ,  $R_3$  will issue a request to and receive an approval from the station-computer A with regard to track  $T_1$ . Assume that train  $R_3$  arrives at A at  $t = 6$  units. At  $t = 5$ , train  $R_2$  will issue a request to A for track  $T_1$  and will be denied. Train  $R_2$  may then reevaluate its routing and assume that it decides to issue a request to A for track  $T_4$ . This track,  $T_4$ , is also occupied until  $t = 7$  and, therefore, the request will be denied. Train  $R_2$  may then revert to its first choice and re-request node A for track  $T_1$ . This process will continue until A grants the request of  $R_2$  at  $t = 6$  units. Then,  $R_2$  will move on  $T_1$  to B at  $t = 8$  units. Similarly, at  $t = 5$  units, train  $R_4$  issues

a request to D for track  $T_3$  and the request is granted. Therefore,  $R_4$  moves toward C via  $T_3$  and will reach C at, say,  $t = 8$  units. Meanwhile, train  $R_1$  arrives at C at  $t = 7$  units and issues a request to D via C for track  $T_3$ . This request is denied. Since  $R_1$  has no alternative, it has to wait at node C until a subsequent request to D for track  $T_3$  is granted at  $t = 8$  units. This process continues until all four trains reach their destinations. It may be observed that the routing decisions are made by the individual train-computers while the track allocations are performed by the station-computers and, as a result, the overall throughput is likely to be significantly higher than that of the traditional algorithm.

Decisions are compute-intensive and as DARYN distributes decisions onto every physical entity of the railway network namely, locomotives and stations, that execute on independent processors, the overall throughput is likely to be significantly higher than the traditional algorithm executing on a uniprocessor. Also, trains communicate with local stations, at best, and with neighboring stations, at worst. Thus, communication time is significantly less than the traditional approach, where every train and track must send information to the fixed dispatcher-computer. Finally, as the number of trains and tracks increase, implying an increase in the size of the system, the computational requirement and the number of computational engines both increase. Thus, it is likely that the total CPU time will increase much more slowly than the corresponding traditional system, implying superb performance for DARYN.

#### 4.2.2 PROOF OF FREEDOM FROM DEADLOCK

DARYN is an asynchronous, distributed approach and it is deadlock-free.

Multiple trains may propagate requests to compete for a single track; all such requests are ultimately routed to a single station-computer that owns the track exclusively. Upon receiving one or more requests, a station-computer must arbitrate, based on the time of arrival, and then generate and propagate a response—approval or denial, within a finite amount of time—to every requesting train. A station-computer's response is computed based on the total number of outstanding requests that it has received and is independent of all other entities in the system. Thus, the decision making by the station-computer is centralized, as opposed to distributed. A station-computer also merely routes messages directed to and from trains and station-computers in finite time. Between the origin and the final destination, every train will continuously issue requests to the stations and await responses. Until a response is received from a station, corresponding to a request, a train will not issue a second request to any station. Every train's decision to propagate a request for a specific track is independent of any other train or station-computer. When a station-computer grants permission to a requesting train, the train then becomes the exclusive owner of the track for the duration of its travel. Thus, there is a complete absence of cyclic dependency and, therefore, the (1) possibility of deadlock is eliminated and (2) the issue of safety, that is, two or more trains may not travel and collide on the same track, is guaranteed. Moreover, for a system with a finite number of trains and the existence of a set of track between the origin and destination for every train, DARYN guarantees that every train will reach its destination in finite time. The proof is

contingent on the assumption that every station-computer handles a finite number of trains.

#### 4.2.3 MODELING DARYN ON A LOOSELY COUPLED PARALLEL PROCESSOR


To analyze its properties in detail, the DARYN algorithm is modeled on a loosely coupled parallel processor, ARMSTRONG [32]. This test bed is appropriate because its architecture—concurrent processors communicating over explicit protocols—closely resembles the geographically distributed and concurrent train- and station-computers that communicate via explicit messages. ARMSTRONG [32] consists of 68 high-performance MC68010 (10 MHz) processors that are connected through high-speed communication links in the topology of a hypercube. Each processor is capable of executing an application program asynchronously and concurrently. Furthermore, processes executing on unique processors may communicate through explicit, high-level communication primitives. In this investigation, for a given railway network, every station-computer is represented through an ARMSTRONG node. The communications channel between two station-computers is modeled through a software protocol established between the corresponding ARMSTRONG nodes. Ideally, one would like every train-computer in the system to be represented also by an ARMSTRONG processor. This would imply that a protocol must be established dynamically between an arriving train-computer and the station-computer where the train has just arrived. In addition, the protocol previously established between the train-computer and the previous station-computer must be deleted because, most probably, it will never be used again. The dynamic allocation and deallocation of protocols, while theoretically supported by ARMSTRONG, is unreliable and, therefore, unusable. In this research, train-computers are conceptually represented through migratory processes. When a train,  $R_i$ , resides at a station A, the corresponding train-computer is implemented through a concurrently executable process, along with the process for the station-computer, on the underlying ARMSTRONG node. When the train moves from A to B, the process on the previous node is disabled and the vital parameters for the train-computer are propagated in the form of a message to B. At B, a new, concurrently, executable process is initiated along with the station-computer process on the underlying ARMSTRONG node. The train-computer process is customized with the parameters of  $R_i$  propagated to B. Thus, a train's functional behavior is executed successively on the ARMSTRONG nodes corresponding to the stations that the train visits until it arrives at its destination. That is, the train's parameters are successively updated and propagated from node to node as the train propagated through the network. A disadvantage of this model is that the train-computers may not be modeled through independent, concurrent, ARMSTRONG processors. Furthermore, a node representing a station has to divide its computational ability between the station- and the train-computer processes. This dual role of the ARMSTRONG node evidently reduces the performance of the DARYN approach. An ideal model of DARYN would implement every train on a concurrent processor of ARMSTRONG.

At every ARMSTRONG node corresponding to a station of the railway network, a complete knowledge of the entire network is stored. When a train arrives at a station, the train-computer process reads this information. Every station-computer is fully aware of the rightful owners of every track in the system. This assumption is based on the idea that the frequency of change of the track layout is very small and any such change information is quickly propagated through the network. A request for a track, issued by a train-computer, is initially propagated to the station-computer where the train currently resides. Then, the message is routed to the station-computer that actually owns and controls the track. Upon receiving two or more requests for a track, a station-computer determines the lowest of all request times and grants approval to the train-computer with the lowest time of request. At this point, all other requests for the track are denied. Where more than one train-computer qualifies, the station-computer issues the approval to only one of them arbitrarily.

Since DARYN is asynchronous, each train-computer maintains its own time counter. Initially, all of the train-computers start with their time counters initialized to zero. As the simulation progresses, the train-computers increment their counters both as they compute train requests and as they move, that is, the time it takes to traverse a track is the distance of that track divided by the speed of the train. Consequently, the advantages of asynchrony are concurrency and lack of complexity associated with synchronizing a large number of trains in a large system.

In ARMSTRONG, the communication process utilizes nonblocking primitives. Once a protocol is established, data may be passed between the nodes as follows. The sending node sends the data to the receiving node via an `ipc_send (information_path, data_1, data_2, ..., data_n)` command. The sender process receives an immediate response from the operating system with regard to the message; where the message is sent successfully, the sending node receives a 1. Otherwise, it receives a 0, implying that the attempt failed. It is the sender process' responsibility to attempt to send the message again in the future. Given that the system is asynchronous, the receiving node may not know the exact time when the sender has sent data. It executes an `ipc_select (information_path)`, once every so often, to check if data has been transmitted along a particular information path. If the result is positive, the receiving node may read that data with an `ipc_receive (information_path, data_1, data_2, ..., data_n)` command. Otherwise, the receiver does not execute an `ipc_receive` command. The `ipc_select` command is also nonblocking in that it will return immediately with either a positive or negative response.

For simplicity, the tracks are laid out in X (horizontal) and Y (vertical) directions with the nodes laid out in a rectangular matrix. This ensures that the track request algorithm for each train remains as simple as possible. Each train calculates its track request by first computing the difference, in the X-Y coordinate system, between its current location and the final destination. If the Y-difference exceeds the X-difference, the train attempts to move in the Y-direction so as to move closer to its final destination. Similarly, where the X-difference is greater than the Y-difference, the train attempts to move in the X-direction. Thus, the train determines the appropriate track request and submits it to the corresponding node processor, which controls the track. Once



**FIGURE 4.4** A potential timing problem in the distributed simulation.

it receives a response—confirmation or denial, the train considers its next move. If a train is confirmed, it moves along the confirmed track. Where the train’s request is denied, it recomputes a second track choice and resubmits it. A train’s second choice consists in its attempt to move in the direction other than its first choice. This process continues until the train is granted approval on a track.

Given that the simulation on the ARMSTRONG system is also asynchronous, two or more messages sent from individual nodes at the same time instant may arrive at their destinations at times that differ greatly. This phenomenon introduces a “timing” problem in the simulation that may be described as follows. Assume that, in a railway system shown in Figure 4.4, two trains are both headed towards the New York station. While train A is traveling on track 2, train B is moving on track 3. Also, assume that upon arrival at New York, both A and B would wish to travel subsequently on track 1. In the real world, the following should occur. Where track 2 is much longer in length than track 3, assuming that both trains travel at the same speed, train B would arrive at New York prior to train A. Thus, the station-computer at New York would most probably grant approval of track 1 to train B. However, in the simulation, the time required to travel is not modeled because the authors intend to measure only the CPU times required by the decision processes. As a result, the following problem may occur. The message containing the parameters for train A arrives at New York prior to train B. That is, given that messages are propagated asynchronously in ARMSTRONG, train A may “electronically” travel faster than B. The station-computer at New York would, therefore, first grant approval to train A, thereby creating an inconsistency. It may be stressed that this problem is purely the result of the hardware limitations of ARMSTRONG.

To address this timing problem, the time required by a train to travel from one station to the next is modeled through a scaled delay. Thus, the actual propagation of a train’s parameters to the subsequent node is physically delayed on ARMSTRONG by this delay amount.

### 4.3 IMPLEMENTATION OF DARYN ON ARMSTRONG

This section presents the data structures and pseudocodes for the different functional units that constitute the implementation.



Each ARMSTRONG node that models a station contains unique information such as the node's unique identifier number, the identifier assigned to the node by the ARMSTRONG operating system, the number of tracks to which it is connected to, the protocol connections for the node, and the delays that are used to simulate real-life computation times. This information is stored in a structure—`node_info`. Each node must also store information on the communication paths such as their unique identifier, the identifier assigned to them by the ARMSTRONG operating system, and two link-lists. The first link-list contains the complete set of input protocols to this node while the second link-list maintains the details of all of the output protocols from this node.

It has been assumed in this implementation that a station-computer controls all tracks to the west and south of itself. It may be noted that, since the track network is rectangular, some station-computers may not control any track at all. For each track that it controls, each node maintains two lists. The first list contains the current reservation requests for that track. This list is processed in every processing cycle and then cleared. The second list stores a single entry corresponding to the train, if any, that currently occupies the track.

Finally, two structures are used by the node for communications. The first structure corresponds to the message between any two nodes. ARMSTRONG requires every message to extend to a finite standard length regardless of the type of message. Since more than one type of message is utilized in the simulation, the structure is necessarily a superset of all of the messages. Conceivably, a message between two nodes may contain a few blank fields depending on the type of the message being propagated. The second structure is utilized by the output buffer for each node. Since nonblocking message communication primitives are utilized, it is possible for message transmissions to fail occasionally. Under these circumstances, the message must be stored in the output buffer to attempt retransmission at a later time. The buffer holds the message until transmission is successful, at which point its copy is removed from the output buffer.

The implementation also utilizes a structure to store the parameters of a train-computer. Given that a train exists only through its parameters, that is passed between the ARMSTRONG nodes, every node maintains an array of the trains currently located at the node.

The host processor, outside of the ARMSTRONG processor, initializes the simulation and loads a copy of the executable program on each of the corresponding ARMSTRONG processors. In the "main" routine, a node first reads-in its numerical identifier that is propagated to it by the host processor. Each node then executes the function "initial" that reads its information from an external input file. The function, `initial`, performs two actions. First, it establishes the appropriate communication paths with the neighboring nodes. To achieve this, it places a function call to the `initialize_protocols` subroutine. This subroutine reads the identities of all of the neighboring nodes from the input file. It then recalls the ARMSTRONG operating system assigned identifier for the neighbor nodes and compares them with its own corresponding identifier. As per assumption, where its identifier is numerically lower than that of the neighbor's, this node sends the path identifier to the neighbor node. Conversely, if its identifier exceeds that of the neighbor's, this node waits to receive the path identifier



```

main ()
{
    initial(train_array); initialize protocol connections, node information,
                        & originating trains
    k_get_uptime(&starttime); initialize starting time of simulation
    simulating = 1;
    while (simulating == 1) {
        if (counter == end_value) stop simulation;
        train_function(train_array); compute train requests for this track
        read_inputs;          scan input protocols for incoming messages
        track_function;      process reservations for all tracks controlled by node
        clear_buffers;      send output messages to other nodes
        traveling_trains; verify whether trains have arrived at the node
    }
}

initial(train_array)
{
    open external input file;
    initialize_protocols;
    read group, node identifier, track coordinates and update structure node_info
    read track direction, identifier, length, nodes connected to track for every
        track and update sub-structure node_info.tracks_connect[ti].
    if (track direction == 3 or 4) update sub-structure node_info.tracks_control.
    read train destination, speed for any train originating at node & update
        the link-list train_array.
}

```

**FIGURE 4.5** The main and initial routines.

from the neighbor node. This process ensures an unambiguous and uniform mode of establishing the protocol connections. Second, the function `initial` reads-in other relevant information with respect to the node such as the neighbor tracks and the trains, if any, that originate at this node. The pseudocode is shown in Figure 4.5.

Following the completion of the initialization routine, each node then initiates a cyclic execution process that terminates at the end of simulation. The cycle, embedded in a while loop, consists of five functions, namely, `read_inputs`, `train_function`, `track_function`, `clear_buffers`, and `traveling_trains`. They are detailed as follows.

1. **Function `read_inputs`:** This function (Figure 4.6) continually checks every input protocol for any incoming messages. It utilizes a for loop to scan through the link-list of input protocols. The first step in the loop is to check for, via an `ipc_select` command, any outstanding data on the protocol. If the result is positive, the message is read through an `ipc_rcv` command. Following execution of the `ipc_rcv` command, the message is sorted with a switch with regard to its message type. Three types of messages are recognized: (i) a track request from a train at another node, (ii) a confirmation or denial of a request sent from a train at this node, and (iii) a train moving into this node. Based on the nature of the message, subfunctions are invoked

```

read_inputs()
{
    while (searching for each input protocol) {
        check for any incoming message with ipc_select;
        if yes, read the message with ipc_receive;
        switch(message_type) {
            message_type = 1: this is a request for a track controlled by this node
                call add_request to include request in the reservation list
            message_type = 2: this is confirmation/denial of a request sent by train at th
                call receive_confirmation to inform the appropriate train
            message_type = 3: a new train moved into this node
                call new_train to put the train on traveling_trains link-lis
        }
    }
}

```

**FIGURE 4.6** The read\_inputs routine.

to perform appropriate tasks of message processing. The subfunctions are described as follows:

- (i) **add\_request:** This subfunction adds a request, generated by a train outside this node, to the appropriate track's reservation list. It achieves this by using a for loop to scan through the tracks that are controlled by this node and adds the new request at the end of the reservation link-list.
- (ii) **receive\_confirmation:** This subfunction updates the confirmation status for a train currently at this node, upon receiving a response from another node. A while loop is utilized to scan through the train link-list and change the appropriate confirmation field.
- (iii) **new\_train:** This subfunction adds a train's parameters to the link-list of trains currently at this node. A while loop is first utilized to determine the end of the train link-list, and then the new train's parameters are attached at this point.

When all new messages have been read-in and processed appropriately, the node waits for a length of time equal to comm\_delay. This refers to the communication delay that has been discussed earlier. Following the delay, the node checks each input protocol again for new messages. If the result is positive, they are sorted by the message type, as discussed earlier.

2. **Function train\_function:** This function (Figure 4.7) simulates the actions of the trains currently located at this node. Utilizing a while loop to scan through the train link-list, each train's confirmation status is reviewed. There are three possible values for the confirmation status, and, corresponding to each of the values, different subfunctions are executed as described subsequently.

- (i) **The value of confirmation is -1.** This implies that the train's second choice has been rejected most recently. Thus, the train needs to recompute its first choice and resubmit it. The subfunction compute\_1 computes a train's first choice track request as described earlier. After the track request is computed, it calls the subfunction send\_request to resubmit the request. If the train has requested a track to the west or south of the node, evidently the current node controls the track and the

```

train_function(train_array)
{
    initialize train_ptr to the dummy entry of train link-list, train_array;
    initialize prev_ptr to the pointer prior to train_ptr;
    while (searching through entries, if any, in the train link-list) {
        if (train has not completed its travel) {
            call check_if_done to verify whether train has reached destination;
            if (train not located at this node) {
                switch (train's confirmation status) {
                    confirmation = -1: train is ready to compute its first choice
                        call compute_1
                        call send_request to submit request for track
                    confirmation = 1: train is ready to compute its second choice
                        call compute_2
                        call send_request to submit request for track
                    confirmation = 2: train has received confirmation for a track
                        call move to simulate propagation of train
                }
            }
        }
        advance train_ptr and prev_ptr to subsequent items in the link-list;
    }
}

```

**FIGURE 4.7** The train\_function routine.

request is filed in the node. Where the train requests a track to the north or east of the node, the request is propagated to the node that controls the requested track. The subfunction send\_request calls the function send\_message\_to\_buffer once send\_request is ready to send a request to another node. Send\_message\_to\_buffer places the track request on the output buffer link-list.

- (ii) **The value of confirmation is 1.** In this case the train's first choice has been denied. Thus, it must compute the second choice and resubmit the request. The subfunction compute\_2 calculates a train's second choice and places on the appropriate track's reservation list by executing the subfunction send\_request. Conceivably, a train may not possess a second choice where the possibility of movement is in a single direction. In such cases the second choice track request is not computed and the train continues to periodically submit its request for the first choice of track.
  - (iii) **The value of confirmation is 2.** This implies that the train has been confirmed for its track request. The subfunction move is executed to propagate the train's parameters to the subsequent node via the send\_message\_to\_buffer function.
3. **Function track\_function:** This function processes the track reservations that are included on the reservation list. First, a check is made to ascertain

whether the track in question is currently occupied by another train. If the result is positive, the track is declared unavailable until the train completes its movement. When the track becomes available, the lowest time of request among the requests in the reservation list is determined and the corresponding train is confirmed. All other requests are denied. Then, the corresponding trains in the reservation list are informed of their confirmation status. A while loop is utilized to scan through the train link-list. Where the train is located at this node, the train link-list is updated. If the train is located at a different node, an appropriate confirmation/denial message is propagated to the corresponding node. Given that the output protocol of the return path is included with the track request message, the function uses this information to send the confirmation/denial message back to the requester. The subfunction `send_message_to_buffer` is utilized to send the message.

4. **Function `clear_buffers`:** This function (Figure 4.8) is responsible for flushing any outstanding, unprocessed messages in the output buffer. A while loop is used to scan each output protocol in the output protocol link-list. Where a message is outstanding at the output buffer, the function attempts to propagate it via the `ipc_send` command. When the message is sent successfully, as indicated by a return code value of zero, the message is removed from the output buffer. Where the propagation fails, it is retained for future attempts to propagate it.
5. **Function `traveling_trains`:** When a train is allocated a track of its choice, it executes a move toward the subsequent station. However, the propagation of the train parameters to the subsequent node must be delayed by an amount that is proportional to the actual length of the track divided by the speed of the train and a scaling factor. This function (Figure 4.9) achieves the desired effect in the following manner. The train's parameters, propagated to the subsequent node in the form of a message, remain unprocessed, that is, the train's subsequent routing is not evaluated, until the train has actually waited at the node for the scaled delay amount. The function first reads the value (e.g.,  $v_1$ ) of the ARMSTRONG node's timer when the train parameters arrive. At a later time, when the timer's current value,  $v_2$ , exceeds  $v_1$  by the scaled delay amount, the train's existence is recorded in the node and it is permitted to recompute its forward route.

```
clear_buffer ()
{
    while (processing every entry in the output buffer) {
        send message on appropriate output protocol;
    }
}
```

**FIGURE 4.8** The `clear_buffer` routine.

```

traveling_trains ()
{
  if a train's parameters are received, record the real time of the processor's clock;
  read the current real time of processor's clock;
  subtract the real time of the train's arrival from the current time to
      obtain the time for which the train has already been delayed;
  compute the scaled time for which the train must be delayed based on the
      length of track and train speed;
  if (the train's actual delay exceeds the required value) {
    search for the end of the link-list for trains at this node;
    append this train to the end of the list;
  }
}

```

**FIGURE 4.9** The `traveling_trains` routine.

DARYN is written in C and is approximately 1500 lines in length. The program is compiled on a SUN 3/60 workstation under the compiler optimization directive “-O4” (level 4) and is executed on the ARMSTRONG system.

#### 4.4 PERFORMANCE OF DARYN

This section details a number of experiments, implemented on the ARMSTRONG system, to evaluate the performance of the DARYN algorithm. A total of four railway networks, consisting of 4, 6, 9, and 12 stations, modeled as ARMSTRONG nodes, respectively, are used in the experiments. The simulation of larger networks is difficult because the ARMSTRONG hardware limits the number of protocols and open file pointers associated with any node at any time instant. The four networks are shown in [Figure 4.10](#). This section also notes the limitations of DARYN.

To evaluate the performance, three measures are proposed. First, the CPU time required for the simulation is compared against that of the traditional centralized approach and speedup factors are computed by varying the size of the network, number of trains, and speed of trains. The speedup factor is defined as the ratio of the CPU time necessary for completion of distributed simulation to that for a corresponding uniprocessor simulation. Given that DARYN is asynchronous the detection of termination is complex. As an approximation, a special train with a very long travel path is included in the system. Its speed is deliberately chosen to ensure that all other trains reach their destinations prior to itself. Thus, the overall CPU time for completion is given by the time the special train takes to complete its journey. The special train originates at a processor and reaches its destination at a different ARMSTRONG node. At the originating node, the value of the node timer is included as a parameter—`start_time`, of the train. Although the nodes are asynchronous, their clock-timers are all reset to zero at initialization and, presumably, the difference in their frequencies is very small. The difference between the `start_time` and the value of the timer of the destination node when the special train arrives there constitutes the required CPU time.




FIGURE 4.10 Four railway networks to evaluate DARYN's performance.

The rationale for the second measure is expressed as follows: if the DARYN algorithm were to exhibit performance scalability, the CPU time required by a train traveling a fixed distance at a fixed speed should not be significantly affected by either the number of trains in the entire system or the size of the network, because any increase in the network size and number of trains must be accompanied by a proportionate increase in the number of computational engines. The

CPU time will be affected somewhat where the number of trains, but not stations, are increased because of the fact that the increased number of train-computers will interact with a fixed number of station-computers. This measure for the distributed simulation is contrasted against that for a corresponding uniprocessor implementation.

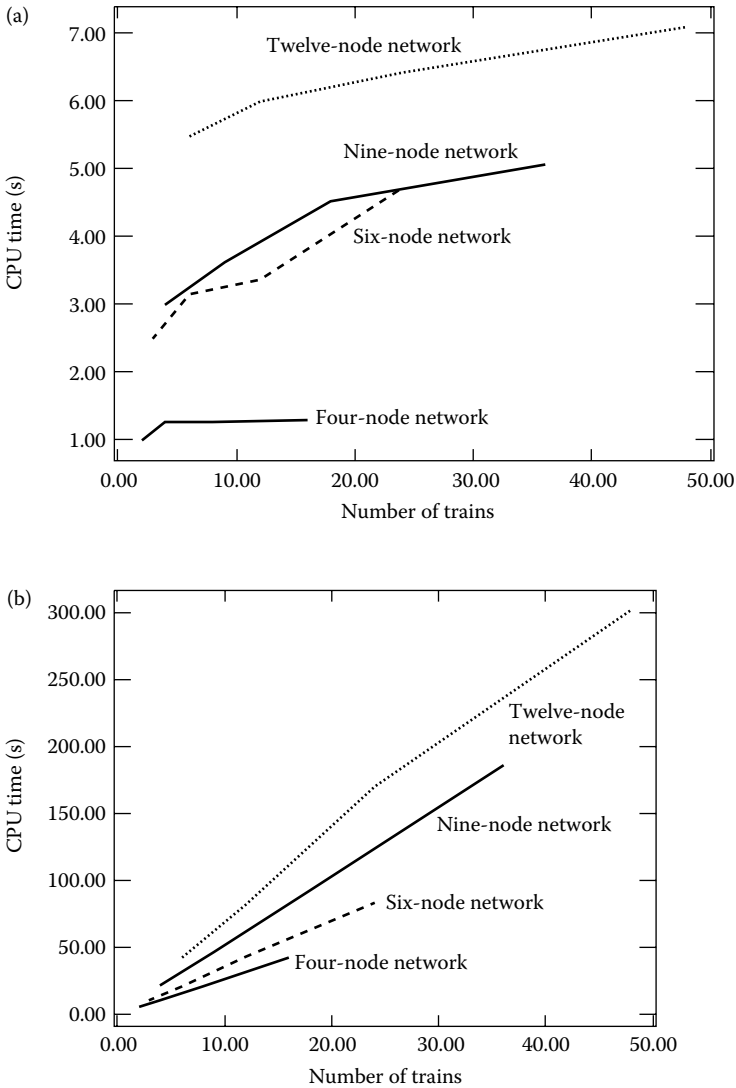
The third measure consists in the measurement of the total “idle time,” that is, the cumulative time for which a train remains stationary at all of the nodes that it encounters while it interacts with the station-computers to receive track allocation. The idle time also includes the times it requires to perform the unsuccessful optimal route computations. This measure for the distributed simulation case is also contrasted against that for a corresponding uniprocessor implementation.

A significant communication delay is required for reading, periodically, the position and speed of every train and the status of every track by the centralized dispatcher in the traditional approach. This is modeled in the uniprocessor implementation through a delay of 0.7 millisecond.

Recall that, in order to address the timing problem in the distributed simulation, the actual travel time was modeled as a delay. In both the distributed and uniprocessor simulations, these delays slowed down the simulation in a consistent manner. While this implied consistency, the CPU times, so obtained, reflect the traveling time in addition to the decision times. That is, the measurements no longer reflect the decision times only. Since the primary goal of this research is to measure the performance of distributed decision making in DARYN, the following was performed. The CPU times for both simulations were obtained by varying the scaled traveling delay and plotted. For both cases, the graphs were observed to be linear, implying that the existence of the traveling delay linearly slowed down both simulations proportionately. From these graphs, the CPU times for the case of traveling delays equal to zero is obtained through extrapolation. For the uniprocessor case, the extrapolated data is observed to be consistent with that from another implementation wherein the traveling delay is not modeled at all.

Furthermore, the uniprocessor implementation is executed on a Sparc 1+ workstation, which is observed to be approximately 8.56 times faster than each processor of the ARMSTRONG [32] system. Thus, all uniprocessor raw data are scaled up by a factor of 8.56 and then compared against the data from the distributed simulation.

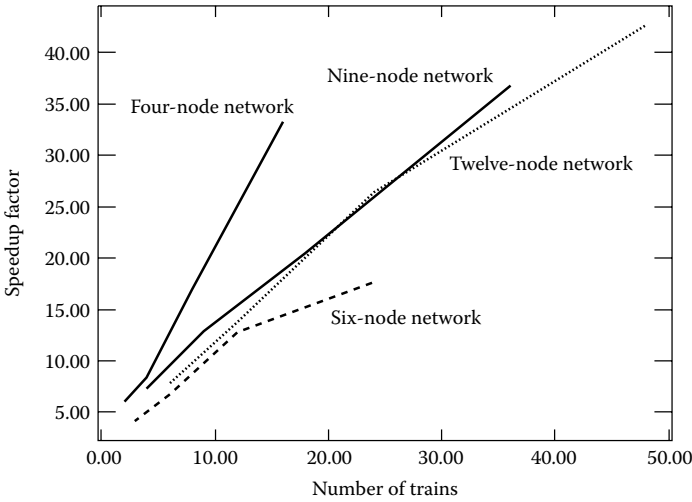
First, a railway network consisting of four nodes and four tracks is considered. The variable, number of trains, is assigned values ranging from 2 through 11, and the corresponding CPU execution times are recorded. In the second, third, and fourth cases, networks consisting of 6 stations and 7 tracks, 9 stations and 10 tracks, and 12 stations and 17 tracks are considered respectively. For each of these cases, the variable, number of trains, is assigned values in the ranges {3, 24}, {4, 36}, and {6, 48} respectively. Also, for all of the experiments, the origin and destination of the trains are chosen to reflect uniformity of distribution. The speed of the trains are assigned at random. The nature of variation of the CPU execution times is observed to be consistent across all the cases. The graphs in [Figure 4.11a](#) present the extrapolated CPU times required for distributed simulation as a function of the number of trains for all of



**FIGURE 4.11** Graphs of CPU times versus number of trains: (a) DARYN algorithm, (b) traditional approach.

the four networks. The graphs in Figure 4.11b describe similar data for the traditional approach. It may be observed that the slopes of the graphs for DARYN are much smaller than those of the traditional approach. The CPU times for Figure 4.11b refer to the scaled data with respect to the Sparc 1+ workstation, that is, they reflect the fact that the original data are multiplied by 8.56 before comparing with the corresponding numbers in Figure 4.11a.



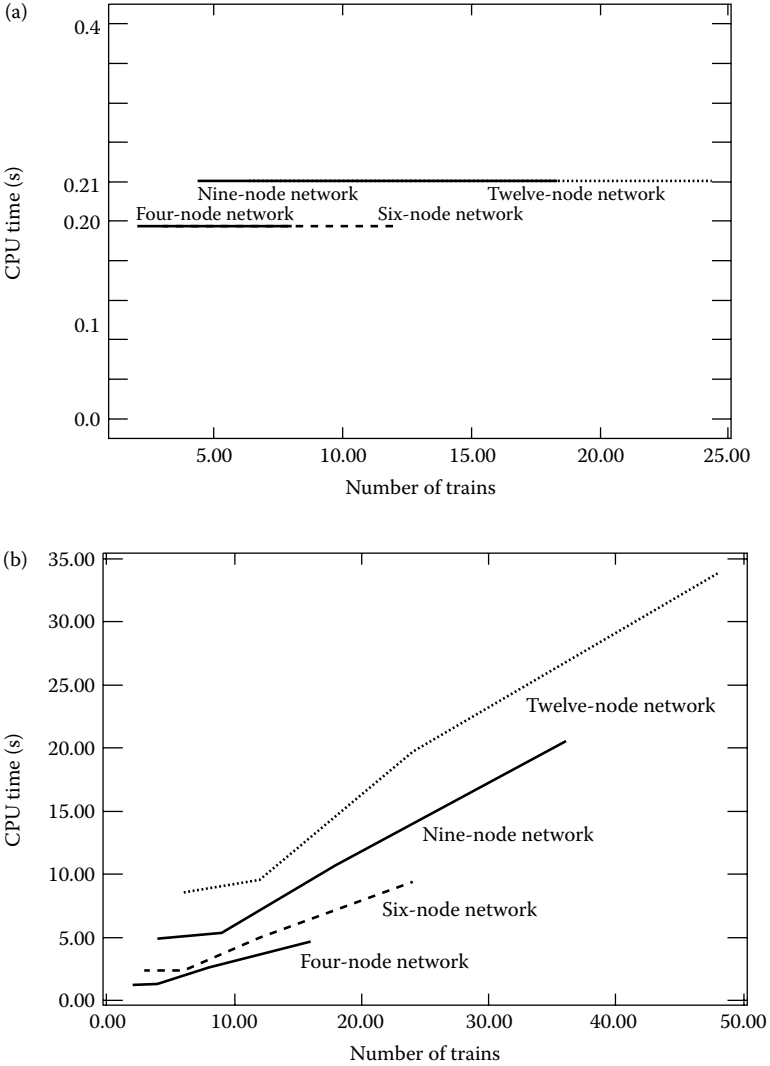


**FIGURE 4.12** Graph of speedup factor versus number of trains for all four railway networks.

The graph in Figure 4.12 presents the speedup factor of the distributed over the traditional approach for varying number of trains and for each of the four networks. The speedup factor is computed as the ratio of the CPU time for uniprocessor simulation to that for distributed simulation. A significant speedup factor of 43 is observed for the case of 48 trains and 12 stations. This speedup, obtained through utilizing only 12 processors, does not violate any of the principles of thermodynamics. It merely reflects the fact that, as the problem size increases, the uniprocessor implementation is increasingly slowed down due to the increasingly larger communication delays.

The graphs in Figure 4.13 present the second measure of performance. The CPU times required by a single train to travel a fixed distance at a fixed speed are noted under different conditions, that is, wherein the number of trains and the network size are both varied. While Figure 4.13a refers to the distributed simulation, Figure 4.13b refers to the uniprocessor approach. It may be observed that while the slopes in Figure 4.13b are significantly large those in Figure 4.13a are, relatively, very small, implying the advantage of distributed decision making over centralized decision making. Furthermore, as the railway network grows in size, as would be expected over time, the performance degradation is relatively minor, implying DARYN’s growth potential.

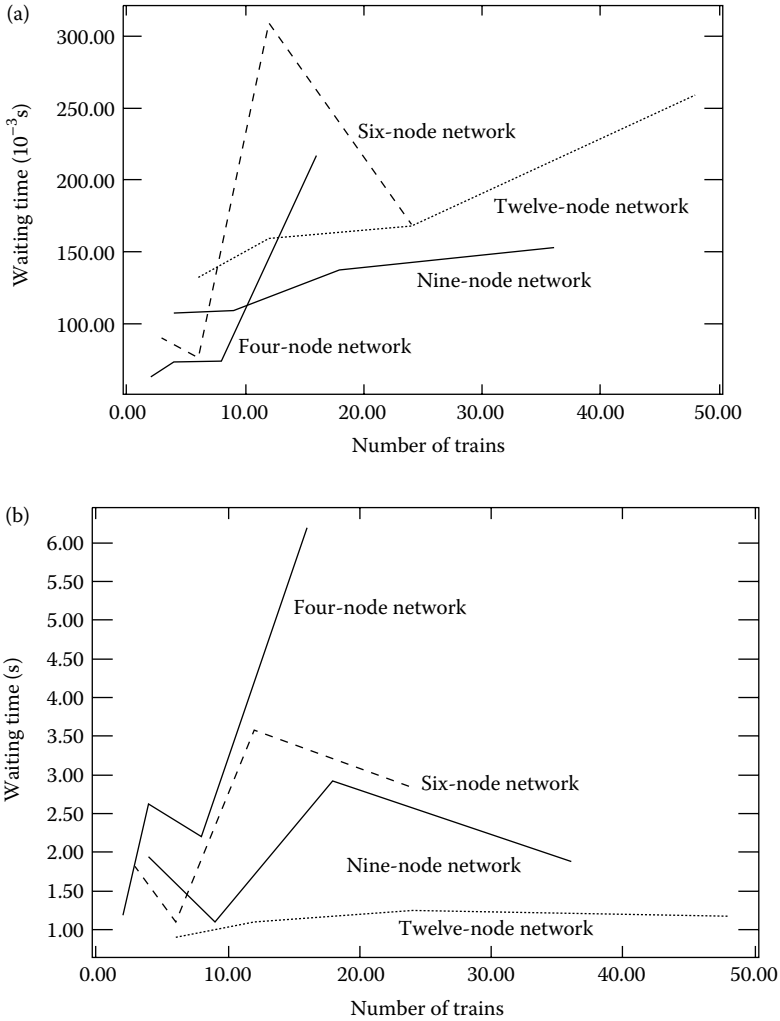
Finally, the cumulative idle time of the special train is plotted for varying number of trains and for all of the four networks. The idle time is determined as the time for which the train is idle at a station while it computes the optimal path, makes decisions, and interacts with the station-computers to receive permission to use its choice track. The idle time is also a measure of the efficiency of DARYN, for it reflects a wastage of valuable resources—trains. The graphs in Figure 4.14a refer to the distributed algorithm and those in Figure 4.14b refer to the traditional approach. Although the graphs do not reveal an easily observable pattern, it is obvious that an idle time in Figure 4.14a is significantly less than the corresponding value in Figure 4.14b.



**FIGURE 4.13** Graphs of CPU times for a train versus number of trains for all four railway networks: (a) DARYN algorithm, (b) traditional approach.

*Limitations of DARYN*

The present implementation of DARYN is limited in that it does not assign priorities to the trains. The station-computers discriminate the trains' requests based solely on the arrival times of the requests and not on the importance of the trains. In the real world, superfast trains or those with critical or perishable cargo may require higher precedence over those carrying nonperishable cargo. The limitation is easily addressed by introducing the notion of priorities in DARYN. DARYN also lacks in-line control



**FIGURE 4.14** Graphs of cumulative idle times versus number of trains for all four railway networks: (a) DARYN algorithm, (b) traditional approach.

for managing crossings and emergencies. To address this problem, additional input signals may be provided to DARYN at the crossings or other critical points in the tracks and emergency information may be issued at the stations. The mechanism of asserting signals at the crossings may be analogous to those that are presently being utilized at the stations. A significant limitation of DARYN is that it fails to include the notions of congestion and bottleneck ahead of its current position, which, in turn, may affect its performance.

---

# 5 RYNSORD: A Novel, Decentralized Algorithm for Railway Networks with Soft Reservation

## 5.1 INTRODUCTION

Railway networks are ubiquitous in today's world, and they continue to play a dominant role in transporting freight and people since 1825, when the first common carrier railroad was introduced. While larger countries such as the United States, Russia, China, India, and the EEC benefit most from extensive and cost-effective railway networks, in many smaller but densely populated countries with large financial resources such as Japan, railway networks contribute significantly to the well being of the national economy by efficiently moving workers and goods. As of 1987, the United States [33] maintains a total of 249,412 miles of railway tracks. It supports a total of 1,249,075,534 locomotive unit miles in one year to carry freight utilizing 5,989,522 loaded cars. For passenger services, the total unit miles stands at 3,782,470 while carrying the gross ton-miles of 1,357,097. In Japan, the East Japan Railway Company (Ryuji Sakamoto, pers. comm., December 1987) carries a total of 16 million passengers each day on 12,000 scheduled trains and 7,500 km of railway track. For efficiency, modularity, and safety, in general, the tracks are divided into individual units, each of which may be controlled exclusively by the system. Thus, a train in propagating from a location A to another location B may travel over several tracks. Given that two or more trains may compete at some time instant for the same track, and that only one train may occupy a track at any time, the principal goal of the railway network management system is to allocate tracks to trains such that (1) collisions are avoided and (2) the resources are optimally utilized.

A detailed analysis of the existing literature in centralized scheduling for railway networks occurs earlier in Chapter 4. In addition, the ASTREE [34] railway traffic management system maintains a distributed database of up-to-date, accurate, and comprehensive representation of route layout and train progress; it uses the information in the database to either automatically make decisions or assist human operators with decisions, relative to route settings and train control. The settings are then downloaded to the wayside equipment and locomotives. Hill, Yu, and Dunn [35] report their effort in modeling electromagnetic interference in railway networks. Ayers [36] presents the use of error-correcting codes to achieve reliable radio-link communication in the advanced train control system (ATCS). Sheikh, Coll, Ayers, and Bailey [37] presents the issue of signal fading in mobile communications. While

Hill [38] presents coding issues to facilitate the communication of train positions efficiently, Shayan, Tho, and Bhargava [39] report of the use of Reed–Solomon Codec to improve the ATCS. The Association of American Railroads (George H. Way Jr., pers. comm., February 1992) notes that distributed algorithms can enhance the efficacy of train scheduling and that several socioeconomic factors, including ownership, track capacity, speed capability, grades, curvatures, clearances, crew districts, and operating agreements, may influence the choice of alternate paths. As noted earlier in Chapter 4, DARYN constitutes a novel, distributed algorithm but is limited in that it employs unit lookahead. That is, at any time instant, it reserves only one track beyond its current position. Consequently, it is unable to utilize congestion information beyond its current position to plan its future route, and this may lead to inefficiency.

This chapter presents RYNSORD, which addresses key limitations of the traditional approaches, described earlier, and marks a significant advancement. RYNSORD studies the concept of lookahead, that is, reserving  $N$  tracks ahead of its current position to improve the utilization of the resources, tracks, and mitigate congestion. It also introduces a new concept, soft reservation, which is characterized by greater flexibility in reservation as opposed to the conventional, hard, reservation technique wherein a reservation request for a specific time instant is either approved or disapproved.

The remainder of this chapter is organized as follows: Section 5.2 presents a detailed description of the RYNSORD approach while Section 5.3 describes the modeling of RYNSORD on an accurate and realistic test bed constituted of a network of 70 SUN Sparc 10 workstations, configured as a loosely coupled parallel processor. Section 5.4 presents key implementation issues. Section 5.5 first reports the performance data from executing a simulation of RYNSORD for realistic railway networks under stochastic input traffic stimulus, and then presents a detailed performance analysis.

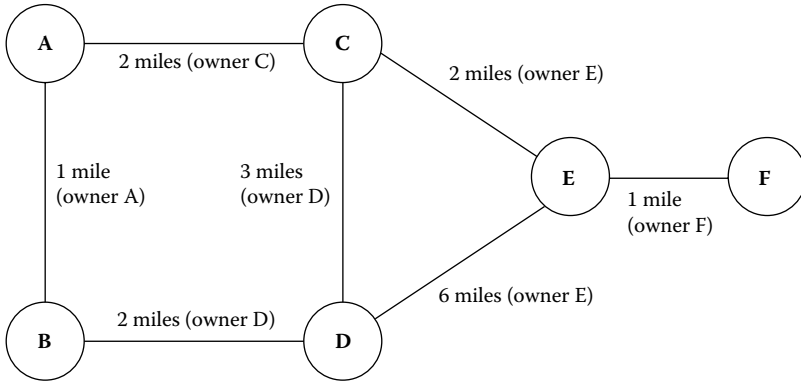
## 5.2 THE RYNSORD APPROACH

The RYNSORD approach for railway networks is novel and defined by the following characteristics: first, it is decentralized in that the overall task of routing all trains through the network is shared by all entities in the system—trains and station nodes. The routing is dynamic, that is, the choice of the successive tracks as a train proceeds toward its final destination takes into account the actual demand of the tracks by other trains in the system. What makes RYNSORD unique among all disciplines including modern communications networks is that every mobile agent, or the train—possesses intelligence, information-gathering facilities, and autonomy to solely and completely determine its own routing. Trains ride on tracks, and safety concerns demand that a train first gains exclusive reservation guarantee from the owner node of a track prior to propagating on it. Conceivably, a train can insist on reserving every track from origin to destination along its chosen route before starting its journey. Such an approach may lock the train to faraway tracks too soon, based on old information, and thereby fail to take advantage of better route choices, which may become available as time progresses. On the contrary, the RYNSORD approach utilizes a lookahead  $N$  wherein every train requests reservations at intervals of  $N$  tracks and a specific

reservation entails the acquisition of approvals of  $N$  subsequent tracks along its route toward its destination before it resumes its travel. Within the reservation process for  $N$  subsequent tracks, RYNSORD proposes a novel concept—*soft reservation*. In the traditional hard reservation, a train issues  $N$  consecutive requests for  $N$  tracks at specific time instances. The owner stations for the corresponding tracks will either approve or disapprove the reservation, depending on whether the respective tracks are free at the requested time instances. Thus, when a train requests a track from time  $t_1$  to  $t_2$  and even if the station notes that the track is occupied up to time  $t_1 + 1$  but free thereafter, it will still refuse approval. Then the train will have to try an alternate track. Assume that the alternate track is a worse solution than if the train had waited idly for 1 time unit and then utilized the original track. If the train had been aware of the knowledge possessed by the station, it could have idled 1 time unit and opted for the better solution. Thus, with regard to reservation, the nodes' behaviors are binary and rigid. In contrast, RYNSORD proposes soft reservation, wherein a train specifies monotonically increasing time instants to the successive station nodes corresponding to the  $N$  subsequent tracks. In turn, a node grants approval at either the requested time instant or the earliest time instant beyond the requested time instants when the track is available. These characteristics are expected to endow RYNSORD with efficiency, robustness, and reduced vulnerability to catastrophic system-wide failures.

In RYNSORD, a railway network is assumed to consist of a set of railroad stations, also termed nodes, that are connected by lengths of railroad tracks. Every station is equipped with a computing engine and communications facilities. For every pair of stations that are connected by a track, there exists a bidirectional communication link between the stations. Every track is bidirectional. Furthermore, every train is equipped with an onboard computing engine and facilities to initiate communication with the computing engine of the corresponding node when it is located at a station. RYNSORD does not require train-to-train communication or train-to-station communication while the train is moving. Every track segment is characterized by its length and the station node that owns it exclusively. A track between nodes X and Y is owned either by X or Y, and the owner is solely responsible for negotiating the use of the track with the competing trains. This characteristic is crucial to guaranteeing safety and collision avoidance in RYNSORD. Consider [Figure 5.1](#), which presents a simple railway network with station nodes A through F that are connected through a partially connected network of track segments. The owners and lengths of the respective tracks are shown in [Figure 5.1](#).

Trains can be asserted into RYNSORD asynchronously, that is, at any arbitrary time, and at any station. Every train is characterized by its originating station, destination station, and its maximum speed. In general, the exact route (the sequence of tracks from origin to final destination) and the consequent arrival time is determined dynamically by RYNSORD. However, where specific intermediate stops are mandated, the overall path is organized into multiple sets of smaller paths, and RYNSORD is applied to each set successively. Thus, if the desired path is  $A \rightarrow C \rightarrow E$ , it is equivalent to first traveling  $A \rightarrow C$ , under RYNSORD, and then from  $C \rightarrow E$ , also under RYNSORD. Excessive use of intermediate stops can lead to poor performance since RYNSORD's strength lies in the dynamic routing of trains to maximize efficiency and resource allocation, and to avoid congestion.



**FIGURE 5.1** An example railway network in RYNSORD.

As indicated earlier, a key concept in RYNSORD is the notion of lookahead, which is defined as the number of track segments that a train negotiates for future use at reservation time. Lookahead reflects how far into the future a train attempts to reserve resources, and these include the subsequent track segments that the train may need to reach its destination, starting with the immediate, next track segment.

Upon entering the RYNSORD system, every train-computer first determines the shortest path between its origin and destination. This is termed the “primary path” and is based on the mileage between the stations. The determination of the primary path does not take congestion into account. A “secondary path” is then determined, whose component tracks are mutually exclusive relative to those of the primary path, with one exception. In a relatively few scenarios, the primary and secondary paths may share one (or more) common track segment if it is the only segment that connects one part to another part of the network. For instance, in Figure 5.1, the  $E \rightarrow F$  link is a necessary component of any path originating or ending at station F and will therefore occur in both the primary and secondary paths.

Next, a train extracts the stations corresponding to the first  $N$  tracks (lookahead =  $N$ ) from both the primary and secondary paths, synthesizes reservation request packets, and initiates them. A reservation request packet consists of a list of successive stations and the expected arrival times at each of the stations. The arrival times are calculated based on the current time, the speed of the train, lengths of the track segments adjoining the stations, and the assumption that trains do not wait at the intermediate stations. That is, the departure time from station  $X$  is identical to the arrival time at station  $X$ . Of course, a train may be subject to waiting at the originating station and other stations where it initiates reservation requests for the subsequent  $N$  tracks. The arrival and departure times determine the time interval for which a track reservation is desired. Thus, for a track segment  $X \rightarrow Y$ , the train must reserve the track for the interval (departure time from  $X$ , arrival time at  $Y$ ).

The train propagates the reservation packet to the first station in the list. If this station is the owner of the first track segment, it will negotiate for reservation for this track. Assume that the train requests reservation for the interval  $(t_1, t_2)$ . If the station

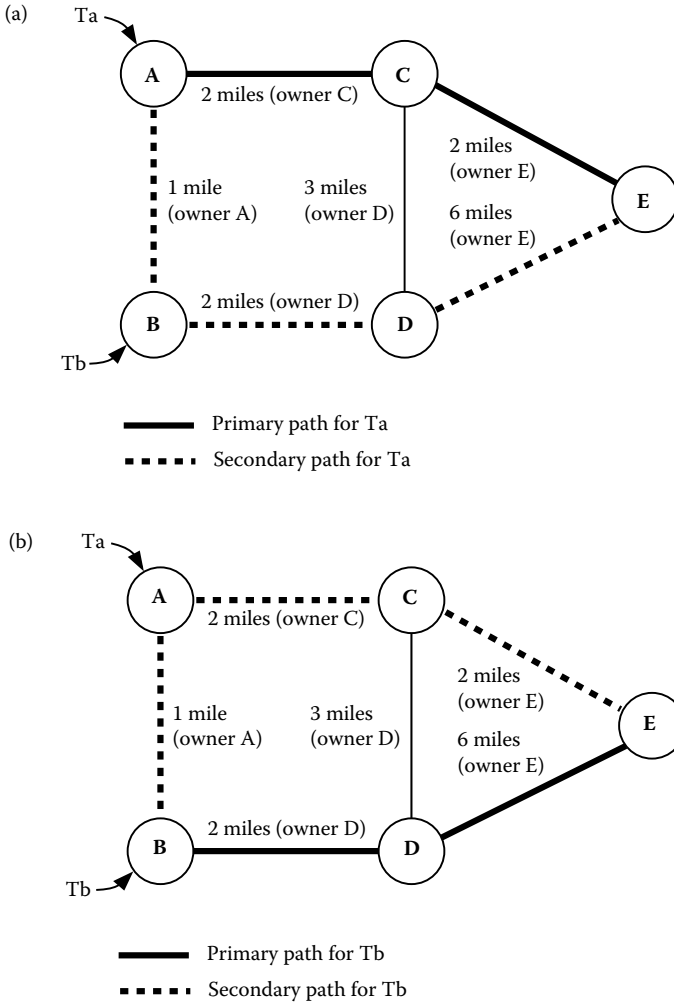
determines that the track is not occupied for this interval, reservation is granted. If, on the contrary, the requested interval is already occupied by another train, clearly reservation cannot be granted for the requested interval. The station then computes the earliest time interval, beyond  $t_2$ , and reserves the track for this new interval, say  $(t_3, t_4)$ . The length of the interval is computed utilizing the length of the track and the train speed. It overwrites the first interval entry in the reservation packet, and the subsequent intervals for the corresponding tracks are also modified. If the first station is not the owner of the first track segment, the reservation packet is forwarded to the second station, which must be the owner of the first track. Following the completion of reservation for the first track segment, the reservation packet is sent to the subsequent station that owns the subsequent track segment. A reservation process, similar to the one described earlier, is initiated, culminating in a reservation time interval for the second track segment. This process continues until reservation intervals are obtained for all  $N$  track segments. The modified reservation packet is then returned to the train located at the station node from where it had initiated the reservation process. This process is executed simultaneously for both the primary and secondary paths.

When a train receives responses to both of its reservation requests along the primary and secondary paths, it may not select as its best choice the route (e.g., R1) that yields the smallest value of time to reach the station at the end of  $N$  subsequent tracks. The reason is that although the primary and secondary paths both lead to the ultimate destination, reaching the end of the  $N$  tracks along route R1 earlier does not automatically guarantee that the train will reach its final destination faster. Therefore, for each of the primary and secondary paths, the train adds the arrival time at the end of the  $N$  tracks to the time of travel from the end of the  $N$  tracks to the final destination along the shortest path. Assume that these times are represented through  $TT_1$  and  $TT_2$  along the primary and secondary paths. The train selects the route that yields the smaller of the  $TT_1$  and  $TT_2$  values. Where  $TT_1 = TT_2$ , the train arbitrarily selects the primary path. Then, the train generates a reservation removal request and propagates it to the stations along the route that is not selected to free the corresponding track reservations.

As a train proceeds from one station to the subsequent station along the  $N$  tracks, it is guaranteed the use of the corresponding tracks in accordance with the reservation times approved earlier. However, should a train arrive at a station earlier than its expected arrival time and if the track is available for a sufficiently long time interval the station may permit the train to proceed immediately. The train, in turn, withdraws the original reservation time interval for the corresponding track segment and modifies its time interval of use of the track. The reason a train, upon arrival at a station node, may find a track available sooner than its requested time interval is because tracks are often freed through reservation removal requests when a train that originally requested reservation decides to select an alternate route. Thus, the previously approved reservation time interval for a track is an upper bound on the travel time for the train. In the event that there are multiple trains competing for a track freed by a train, the train that has been waiting the longest at the station is given the highest preference.

To understand the operation of RYNSORD, consider the railway network in Figures 5.2a and 5.2b that is identical to that in Figure 5.1 except that node F is missing. Assume that two trains, Ta and Tb, are asserted into the system at the same





**FIGURE 5.2** Computation of primary and secondary paths in RYNSORD: (a) for train Ta, (b) for train Tb.

time,  $t = 0$ , at nodes A and B respectively. Both Ta and Tb are destined for station E. Figure 5.2a describes the computations of the primary and secondary paths for Ta, from the origin A to the destination E. Figure 5.2b describes the computations for the primary and secondary paths for Tb, from origin B to destination E. In both Figures 5.2a and 5.2b, the solid and dotted lines represent the primary and secondary paths respectively. Assume that the lookahead  $N = 2$ .

Assuming the value of lookahead,  $N = 2$ , trains Ta and Tb extract the stations relative to  $N = 2$  tracks from both primary and secondary paths. For this example, assume that the primary path for Tb is selected based on the number of tracks, not the

mileage from the source to the destination. The stations for the primary and secondary paths are:

Train Ta:

Primary path station list: A->C->E  
 Secondary path station list: A->B->D

Train Tb:

Primary path station list: B->D->E  
 Secondary path station list: B->A->C

Assuming the speeds of the trains at 1 mile/min, the reservation request packets generated and initiated by Ta and Tb are:

Train Ta:

Primary path:

arrival at A at time 0  
 departure from A at time 0  
 arrival at C at time 2  
 departure from C at time 2  
 arrival at E at time 4  
 [no departure since E is the final destination]

Secondary path:

arrival at A at time 0  
 departure from A at time 0  
 arrival at B at time 1  
 departure from B at time 1  
 arrival at D at time 3  
 [no departure since D is the last station in the station list]

Train Tb:

Primary path:

arrival at B at time 0  
 departure from B at time 0  
 arrival at D at time 2  
 departure from D at time 2  
 arrival at E at time 8  
 [no departure since E is the final destination]

Secondary path:

arrival at B at time 0  
 departure from B at time 0  
 arrival at A at time 1  
 departure from A at time 1  
 arrival at C at time 3  
 [no departure since D is the last station in the station list]

Figures 5.3a through 5.3d describe the operation of RYNSORD relative to the reservation packet propagation by the trains and their processing by the respective stations. In Figure 5.3a, Ra1 and Ra2 represent the reservation packets propagated by Ta along the primary and secondary paths. Rb1 and Rb2 represent the corresponding packets for Tb. In Figure 5.3a, Ra2 requests station A to reserve the track  $A \rightarrow B$  for time (0,1). The request is approved successfully since the track  $A \rightarrow B$  is free for the time interval (0,1). Since station A does not own track  $A \rightarrow C$ , Ra1 cannot utilize station A to accomplish its goal. Neither Rb1 nor Rb2 are able to utilize station B since the latter neither owns track  $B \rightarrow D$  nor  $B \rightarrow A$ .

Figure 5.3b represents the propagation of the reservation packets to the subsequent stations. Here, Ra1 and Rb1 successfully reserve the tracks  $A \rightarrow C$  and  $B \rightarrow D$  for time intervals (0,1) and (0,1) respectively. Ra2 fails to accomplish anything since B is not the owner of track  $B \rightarrow D$ . When Rb2 attempts to reserve track  $B \rightarrow A$  for the time interval (0,1), it fails since train Ta has already reserved that interval. Therefore, station A reserves the next available time interval (1,2) for train Tb. Train Tb updates its reservation packet for the secondary path as shown through the compact representation— $[B@0/1][A@2/X][C@X/X]$ , which implies that train Tb waits at station B from time 0 to 1, then proceeds to station A at time 2, and then departs from A at time 2 to arrive at station C at time 4. Train Tb is restricted from reserving tracks beyond station C by the lookahead value of 2, and this is reflected by the subfield “[C@4/X],” where X implies unknown. Each subfield of the compact reservation packet represents [station name@arrival time/departure time (X implies unknown)].

Figure 5.3c represents the subsequent propagation of the reservation packets. Ra1 and Rb1 successfully reserve the tracks  $C \rightarrow E$  and  $D \rightarrow E$  for time intervals (2,4) and (2,8) respectively. Ra2 fails to reserve the time interval (1,3) on track  $B \rightarrow D$  as train Tb has already reserved the interval (0,2). Train Tb is allowed reservation for the time interval (2,4), and its compact reservation representation is  $[A@0/0][B@1/2][D@4/X]$ . Rb2 succeeds in reserving track  $A \rightarrow C$  for the time interval (2,4).

All of the reservation packets Ra1 through Rb2 have successfully reserved the last track under lookahead = 2 and are returned to the respective trains Ta and Tb at stations A and B respectively. At node A, train Ta notes that the total time to reach the destination E through the primary path is 4. The secondary path requires 4 time units to reach D, and the extra travel time to destination E will demand at least 6 time units, implying a total travel time of  $4 + 6 = 10$  time units. Clearly, train Ta selects the primary path, that is,  $A \rightarrow C \rightarrow E$ , as the best choice and then propagates a reservation removal request to the stations contained in Ra2. For train Tb, the primary path requires 8 time units to reach the destination E. The secondary path requires 4 time units to reach C and the extra travel time to destination E will demand at least 2 time units, implying a total travel time of  $4 + 2 = 6$  time units. Therefore, train Tb selects the secondary path, that is,  $B \rightarrow A \rightarrow C$  and propagates a reservation removal request to free reservations that it had earlier acquired along the primary path.

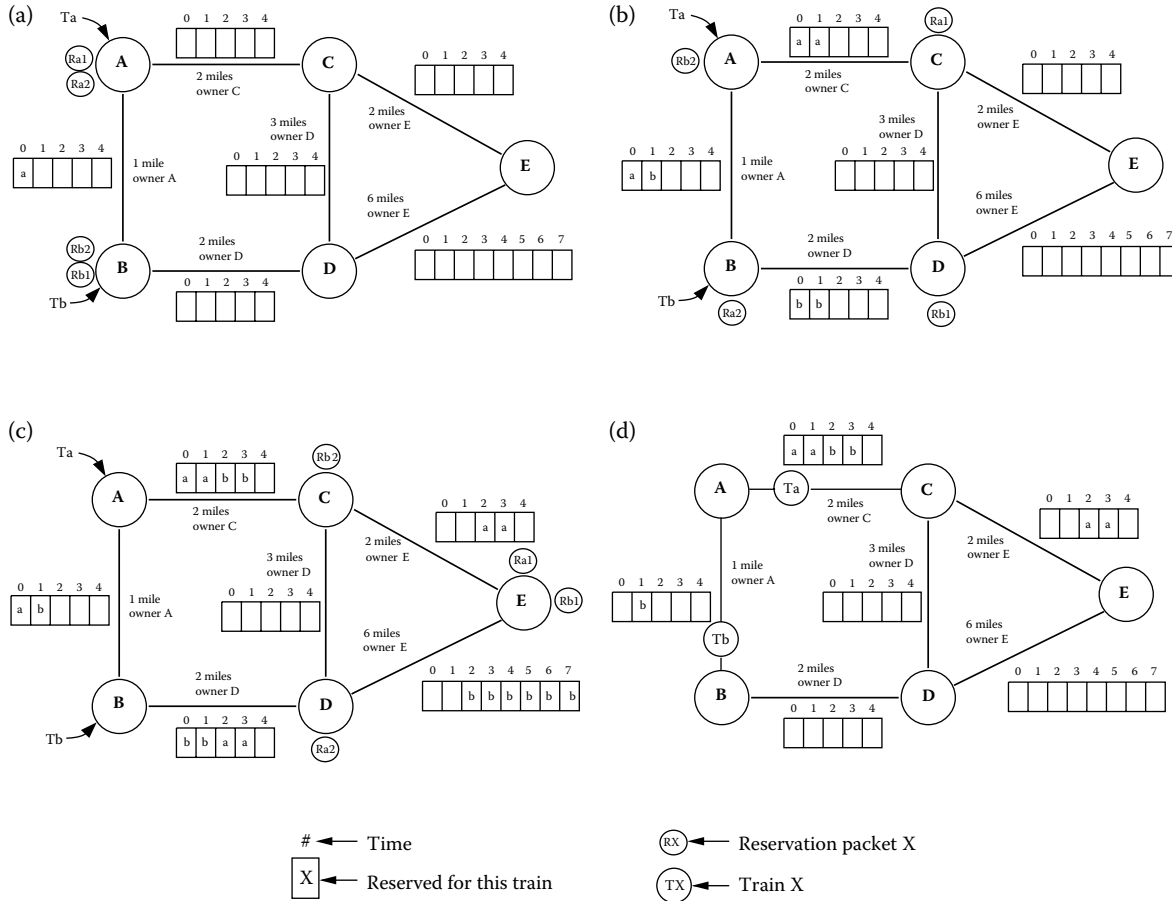


FIGURE 5.3 Initiation and progress of reservation processing.

Figure 5.3d represents RYNSORD when the trains  $T_a$  and  $T_b$  have started to travel and the reservation removal requests have been processed. It may be noted that following the removal of the reservation for train  $T_a$  on track  $A \rightarrow B$ , conceivably train  $T_b$  may be permitted to travel earlier than its reserved time interval of (1,2).

### 5.3 MODELING RYNSORD ON AN ACCURATE, REALISTIC, AND PARALLEL PROCESSING TEST BED

The key contribution of RYNSORD is in distributing the overall task of routing all trains through the network among all the entities in the system—trains and station nodes. Thus, in reality, for a given railway network with  $N$  trains and  $S$  stations, the total number of coordinating computing engines is  $(N + S)$ . To understand its performance and its dependence on different factors, RYNSORD is first modeled and then simulated on a parallel processing test bed that is constituted of a network of workstations configured as a loosely coupled parallel processor. The simulation coupled with the test bed virtually resembles a real implementation with one exception. To facilitate the simulation of a realistic system, that is, with a reasonable number of trains, while every station node is represented by a workstation, the trains are modeled as tasks and executed by the workstations underlying the stations. When a train is located at a host station, its computations are performed by the underlying workstation and its communications with other stations are also carried out through this station. When a train travels from the current station (say  $A$ ) to another station (say  $B$ ), the corresponding train-task in the underlying workstation for  $A$  is encapsulated through a message, propagated to  $B$ , and remanifested as a train-task in the underlying workstation at  $B$ . Thus, trains move in the simulation at electronic speeds instead of their physical speeds and a train's computation and communication subtasks are executed on the host station's underlying workstation.

While trains propagate at approximately 120 mph, the underlying, fast, computing engines of the test bed enable the simulation to execute many times faster than reality. This, in turn, facilitates the rapid performance evaluation of RYNSORD for different values of the parameters. The basic unit of time in the simulation is termed a timestep and it defines the finest resolution of train movements. Thus, while the smallest distance traveled by any train is limited to that within the interval of 1 timestep, a train must also wait at a station for at least 1 timestep if it has not received the necessary reservations to commence travel. In the current implementation of RYNSORD, the timestep value is set to 1 minute of actual operation. The principal reasons for this choice are that (1) the distance traveled by the fastest train, namely 2 miles, is significantly smaller than the shortest track of length 50 miles and (2) relative to processing a reservation for  $N$  tracks, all necessary computing and electronic communication between stations and trains may be accomplished within 1 minute. While a train requires a period of minutes to travel a single track, a message propagation and computing function only requires of the order of 10 milliseconds. RYNSORD permits trains to be introduced into the system asynchronously, that is at irregular intervals of time. In addition, the trains themselves are autonomous and, therefore,

their decisions are executed asynchronously with respect to each other. Furthermore, the test bed consists of heterogeneous workstations that may have differing clock speeds. Therefore, for consistency and accuracy of the propagation of the trains in the system, RYNSORD requires that the timestep values of every station node and train be synchronized. This guarantees that if two trains,  $T_a$  and  $T_b$ , reach their destination, station E, at actual times 12:01 PM and 12:05 PM, for example, in the corresponding RYNSORD model,  $T_a$  must arrive at E prior to  $T_b$ , despite differing processor and communications link speeds. Synchronization is achieved in RYNSORD through an asynchronous, distributed, discrete event simulation technique utilizing null messages [40,41] and is not detailed here.

The previously stated assumption that all message communications and decision processes relative to a reservation request must be completed within a timestep implies the following. If a train  $T_a$ , at a station A, initiates a reservation packet at timestep  $t_1$  and propagates it to other appropriate stations (X, Y, Z, . . .), the reservation packet must be processed at the appropriate stations and returned to  $T_a$  at A prior to advancing the timestep value at A and every one of the workstations underlying the stations (X, Y, Z, . . .) to timestep value  $t_1 + 1$ . To achieve this objective, RYNSORD employs a special, synchronizing node that is connected to all station nodes. It monitors whether all necessary communications and responses corresponding to all reservations that are launched out of the stations, if any, are completed before permitting the station nodes to increment their timestep values. The special, synchronizing node in RYNSORD is an artifact of the parallel simulation on the test bed and has no counterpart in reality. In an actual railway network, electronic communication and computations will require approximately 10–100 milliseconds within which the fastest train will have traveled a mere 18 feet of track. Thus, for all practical purposes, communications and computations are instantaneous.

The basic functionalities of every station node and the special, synchronizing node are encapsulated in pseudocode, as shown in Figures 5.4 and 5.5 respectively.

RYNSORD requires four types of communications messages—first to represent

```

while simulation is not finished {
  send out reservation requests
  while not done {
    process incoming trains
    process incoming reservation requests
    process incoming reservation responses
    if received responses to all reservation requests
      send done to synchronization node
    if received update from synchronization node
      set done to true
  }
  update internal time
}

```

**FIGURE 5.4** Functionality of a station node in pseudocode.

```

while simulation is not finished {
  while not done {
    if received done from station node
      increment count
    if count is equal to number of stations
      set done to true
  }
  send update to all stations
  update internal time
}

```

**FIGURE 5.5** Functionality of the special, synchronizing node.

“reservation packets” that are initiated and launched by trains, second to model the encapsulation of and propagation of trains from one station to the subsequent station, third to allow a train to negotiate with the owner station for earlier than scheduled travel on the subsequent track if it arrives early at a station at the head of the track, and fourth to allow a station grant permission to a train to travel on a track immediately. Reservation packets are assumed to propagate without any delay, that is, within the same timestep. Once received at a station, they are processed immediately and within the same timestep. The final, approved reservation packets are also returned to the originating trains within the same timestep. At the originating station, initially a train lacks a reservation packet. It creates and then launches the reservation packet. Upon receiving the approved reservation packets, the train selects one of them as its best choice, which thereafter becomes its reservation packet. The reservation packet, however, is only good for  $N$  subsequent tracks, and the train will need to repeat the process until it reaches its final destination. A reservation packet is characterized by five fields that are enumerated and explained as follows.

*Reservation packet:*

1. **Station list:** The complete station list, including the station id, arrival time, and departure time, for each station.
2. **Status:** The status of the reservation—(1) **RESERVING** if the reservation packet is traveling forward through its station list attempting to make reservations, (2) **REMOVING** if the reservation removal packet is propagating forward while releasing the previously granted reservations, or (3) **ACCEPTED** if the reservation packet is returning to the originating train.
3. **Train id:** The unique identification number for the train. The  $id = (\text{originating station} * 100,000) + \text{time at which the train is introduced into the system}$ .
4. **Train speed:** The speed information is necessary when the station must modify the reservation time for the train since the original request cannot be satisfied. The speed information is used to compute the travel time over the track.
5. **Reservation id:** The unique identification number for the reservation.

Since a train may require a substantial amount of travel time from one station to the subsequent station, the “train packet” is not assumed to propagate instantaneously. In fact, every train packet is labeled with a timestamp value that represents the timestep at which it is expected to arrive at the destination station. Upon receiving a train packet, a station node stores it in a buffer until the timestamp value of the packet equals the station’s own timestep value. Then the train is assumed to have arrived at the station node and further processing is initiated. A train packet consists of six fields that are enumerated and detailed subsequently.

*Train packet:*

1. **Timestamp:** The expected arrival time of the train at the receiving station.
2. **Train id:** The unique identification number for the train. The  $id = (\text{originating station} * 100,000) + \text{time at which the train is introduced into the system}$ .
3. **Origin:** The originating station of the train.
4. **Destination:** The final destination station of the train.
5. **Path:** A sequential list of the tracks traveled by the train for the purposes of data collection and analysis.
6. **Reservation:** The reservation packet associated with this train.

Since a train may request cancellation of previously approved reservations for tracks along a path when it decides to select an alternate path, conceivably, a train upon arrival at a station may find its subsequent track unoccupied. For efficient use of tracks, the train must be allowed to travel along the track, if possible. To achieve this objective, when a train arrives at a station before its scheduled departure time from that node, it generates and propagates a “waiting packet” to the station that owns the track. Upon receiving the waiting packet, the corresponding owner station queues the train. At every timestep, the station examines whether the track is free and notes the number of timesteps (say  $Q$ ) for which the track is unreserved. The station then selects from the queue a train, if any, that may successfully complete the travel within  $Q$  timesteps and that has been waiting the longest. The station sends a “permission packet” to the train, allowing it to use the track immediately and removes the corresponding entry from the queue. The waiting packet consists of four fields, as shown below, and the permission packet contains a single field.

*Waiting packet:*

1. **Train id:** The unique identification number for the train. The  $id = (\text{originating station} * 100,000) + \text{time at which the train is introduced into the system}$ .
2. **Train speed:** The train’s speed, which is needed to calculate the travel time over the subsequent track.
3. **Wait start time:** The timestep at which the train arrives at the station earlier than its scheduled departure time and is queued.



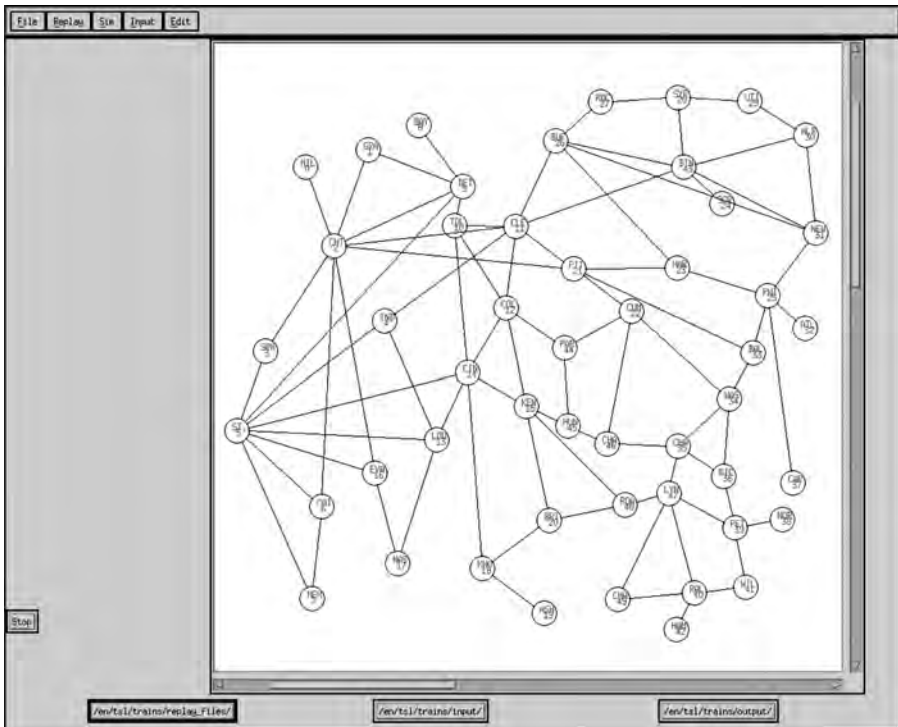
- 4. **Location:** The station, at the head or tail of a track, where a train is waiting. This information is used by the station that owns the track to direct the permission packet, if and when necessary.

*Permission packet:*

- 1. **Train id:** The unique identification number for the train. The  $id = (\text{originating station} * 100,000) + \text{time at which the train is introduced into the system}$ .

To facilitate understanding the distributed, dynamic routing of trains in RYNSORD and the impact of different parameters on the performance of RYNSORD, a visual display of the operation of RYNSORD is achieved through the use of a graphical front end. The graphics supports the following characteristics:

- Developing and editing a railway network.
- Viewing a replay of a simulation run.
- Monitoring and interacting with the simulation, at runtime.
- Viewing statistical information related to the input traffic, results, and runtime performance of a simulation run.

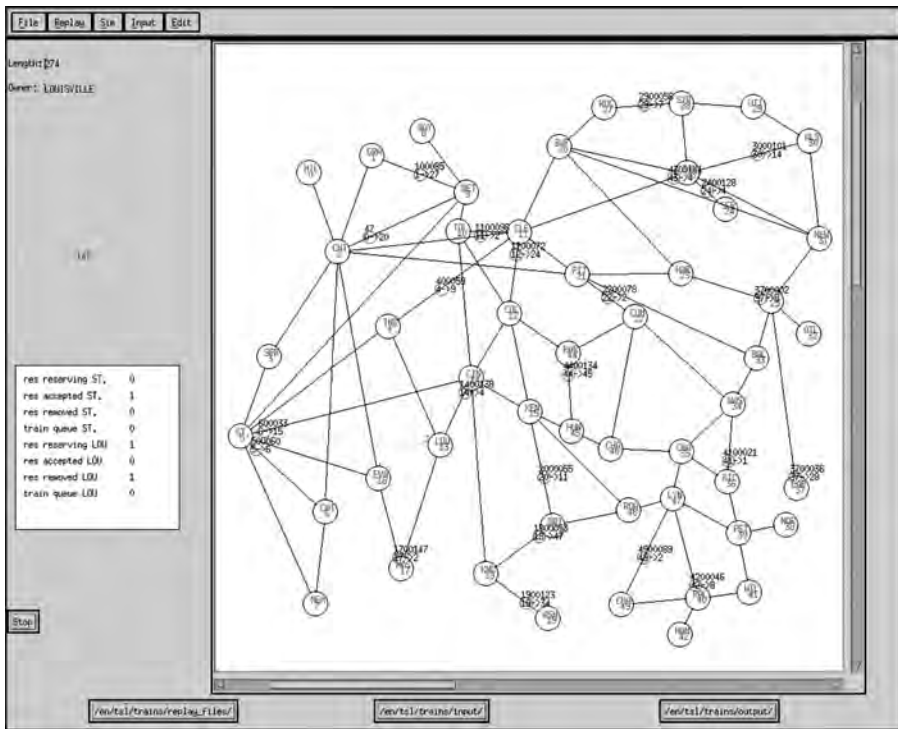


**FIGURE 5.6** A screen shot of the graphical interface.

The runtime display also shows the following parameters:

- The location of train.
- The cumulative number of reservations processed at each station.
- The number of trains waiting at each station.
- The cumulative number of reservations propagated along every segment, categorized by type—Reserving, Removing, or Accepted.
- The cumulative number of trains that have propagated over each track.

Figure 5.6 presents a screen shot of the graphical interface that displays the 50-station railway network detailed in Figure 5.8. In Figure 5.6, each station is labeled by the first three characters of its name and its unique identification number. Stations and links are easily added or deleted directly through this interface and the graphical program will reconfigure RYNSORD, automatically and correctly, to execute the simulation accurately. Figure 5.7 presents a screen shot of an actual simulation run. The trains are described through circles and are located on top of the tracks on which they are traveling. They are identified by their respective identifiers, origins, and destinations. During an actual simulation run, a user may interactively communicate with



**FIGURE 5.7** A screen shot of an actual simulation run.

any of the stations, retrieve any desired data structure and information, and generate and introduce trains into the system at any timestep and at any station.

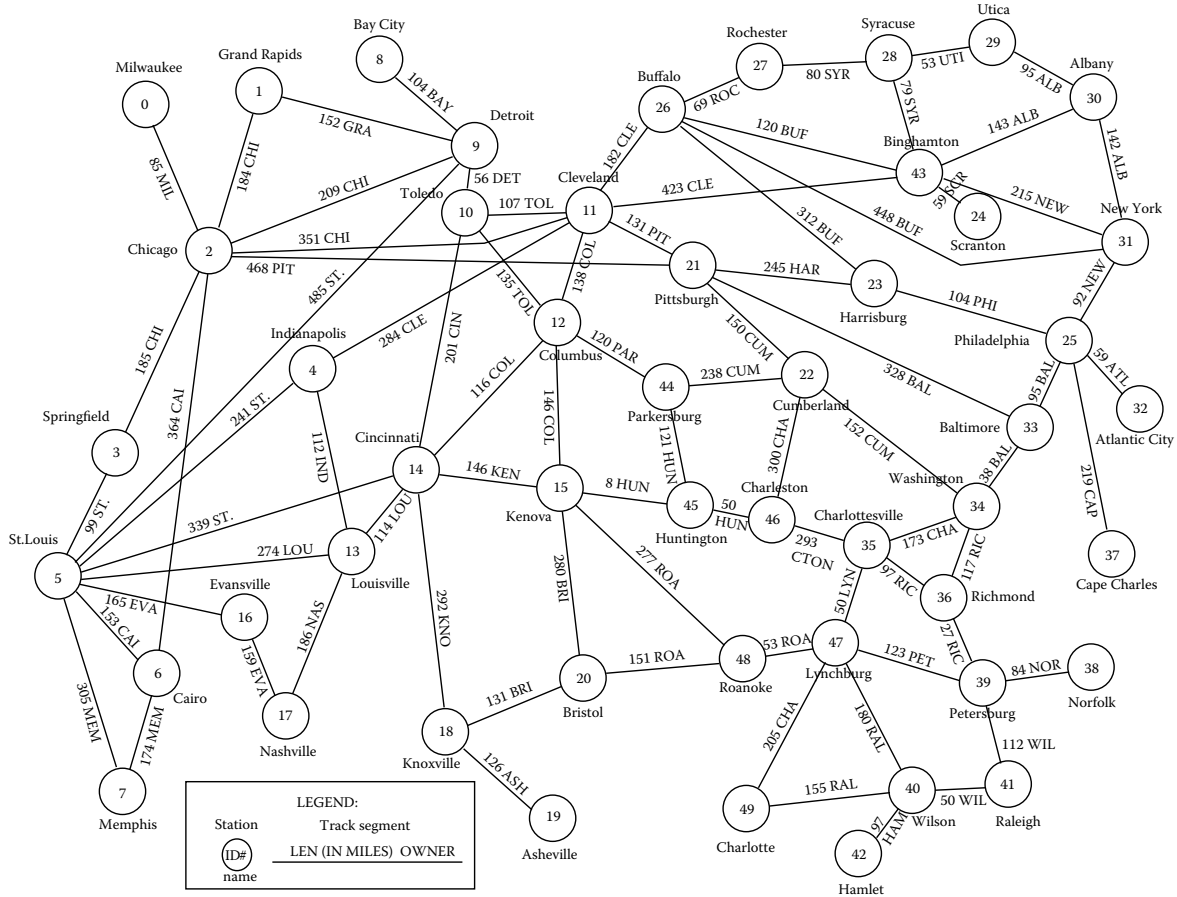
## 5.4 IMPLEMENTATION ISSUES

The RYNSORD model and simulator is written in C and designed to execute on a heterogeneous network of Unix-based workstations, connected through a 10 Mbit/s Ethernet and configured as a loosely coupled parallel processor. The workstation mix includes SUN Sparc 1, Sparc 2, and Sparc 10 under SUN OS 4.1.2 and SUN Solaris 5.3 operating systems and Intel 486DX2/66 and Pentium under the freely available and copyrighted Linux operating system. Station nodes and trains are modeled through processes, and they communicate through TCP/IP. The code segment for every station including the trains located at it is approximately 1700 lines of C code while the networking code is approximately 1000 lines of C code. The simulator is compiled by the public domain GNU C compiler, gcc, and executed at a low priority in the background, utilizing the “nice” utility. It may be noted that the workstations may be executing primary jobs for the users at the consoles. With 50 SUN Sparc 10 workstations executing concurrently, the average execution time of a single simulation experiment is approximately 2 hours of wall clock time.

RYNSORD is an application program that is built on top of the transport payer, TCP/IP. By definition, the layers in the ISO-OSI model starting with the session and beyond are responsible for any required data conversion while communicating between two or more machines. Since Intel 80X86 machines employ Little Endian model while the SUN Sparc machines utilize the Big Endian model [42], the necessary conversion of data in the heterogeneous network of workstations is achieved through the use of n-to-h l (network to host long integer), n-to-h s (network to host short integer), h-to-n l (host to network long integer), and h-to-n s (host to network short integer) utilities [43].

In this chapter, a subset of the eastern United States railroad network is selected, based on the existing primary railroads in the eastern United States, as shown in the 1994 *Rand McNally Commercial Atlas*. A few additional tracks are added to represent a few secondary railroad segments. Figure 5.8 presents the representative railway network that consists of 50 major stations, 84 track segments, and a total length of 14,469 miles of track. A model of the network in Figure 5.8 is developed in RYNSORD and the simulation is executed on a network of 50 workstations, with one station being executed on a workstation.

To obtain representative performance results, a number of experiments are executed with input trains generated stochastically and asserted into RYNSORD. Guidance relative to the choice of density of train traffic, that is, the number of trains introduced into the system per unit timestep is obtained from the actual number of freight trains per 365 days in a year that utilize the tracks of the eastern United States railroad. For the experiments in this chapter, three train traffic densities are selected—low, medium, and high. For low, medium, and high traffic densities, the probabilities that a train is generated within a timestep, that is, 1 minute of real time, are set at 0.075, 0.15, and 0.30 respectively. For every train originating at a station, train speeds



**FIGURE 5.8** A representative railway network: A 50-station subset of the eastern United States railroad network.

**TABLE 5.1**  
**Input Train Traffic Parameters**

<b>Input Traffic Density</b>	<b>Cumulative Trains Introduced in RYNSORD</b>	<b>Estimated Cumulative Distance Traveled by All Trains (miles)</b>
low	484	288,860
medium	869	497,658
high	1,772	1,033,123

are generated stochastically, and they range from 60 mph to 100 mph. The final destination is also generated stochastically by assigning equal weight to every station, except the originating station, and selecting a station at random. Geographic proximity plays no part in the selection process. Since major stations, corresponding to major urban cities, are more likely to encounter high traffic densities, a set of nine “high traffic” stations are selected from Figure 5.8. They include Chicago, Detroit, St. Louis, Philadelphia, New York, Washington, Pittsburgh, Columbus, and Cincinnati. For the stations corresponding to these cities, the input train traffic densities are assumed to be doubled. Also, during the process of selecting final destinations of trains, these cities are assigned twice the weight of other stations.

The representative railway network is simulated in excess of 150 times, under different scenarios and for different parameters. Every simulation is executed for 10,080 timesteps that corresponds to one week of real time operation. As indicated earlier, while a typical simulation experiment executes for approximately 2 hours of wall clock time, the longest running of the 50 workstations often executes for 7 hours. Input trains are introduced throughout every simulation run at a constant and uniform rate that is set at the start of the simulation. Table 5.1 presents the cumulative number of trains introduced into the system and the estimated cumulative miles traveled by the trains, for each of the low, medium, and high input train densities. The estimate is based on the assumption that every train actually travels along the shortest path from the origin to destination, which may be not be true in every case.

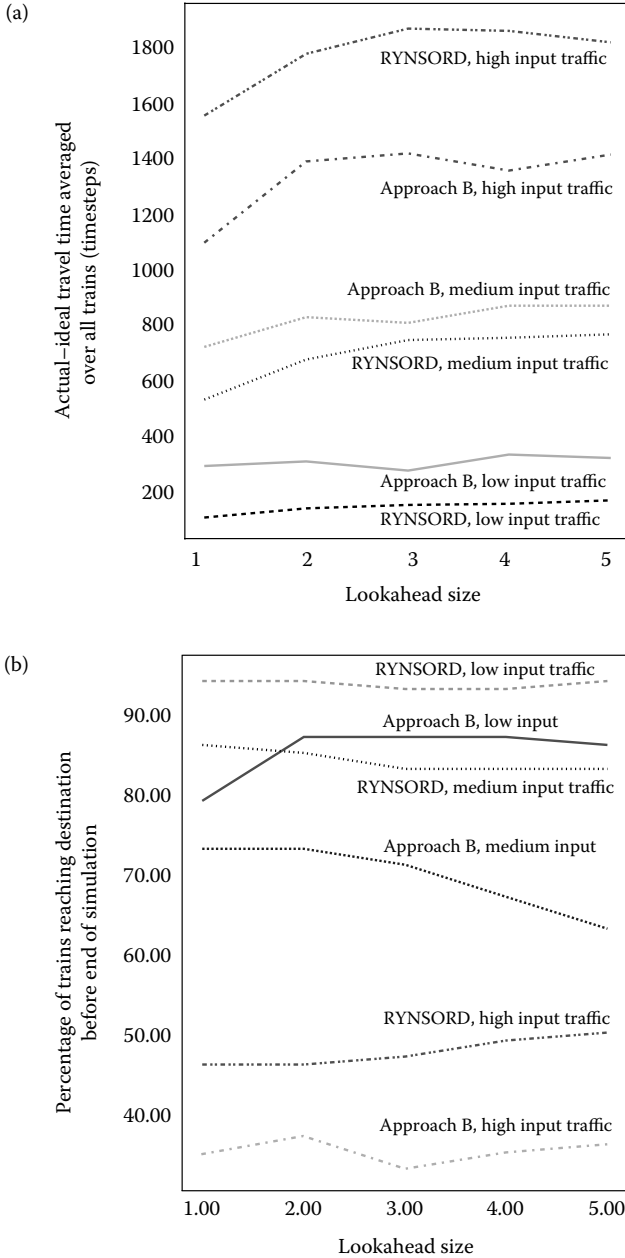
## 5.5 SIMULATION DATA AND PERFORMANCE ANALYSIS

To understand the performance of RYNSORD, first the independent parameters and key performance measures are identified and a number of simulation experiments are executed on the realistic railway network, presented in Figure 5.8. The independent parameters include the (1) number of trains asserted into the system, (2) the density of trains, that is, the frequency with which the trains are input into the system, and (3) the lookahead utilized. To evaluate the role of soft reservations objectively, this chapter also implements a competing distributed routing algorithm, referred to as approach B. Approach B is similar to RYNSORD in all respects except that it employs the traditional, hard reservations policy. A train first sends out hard reservation

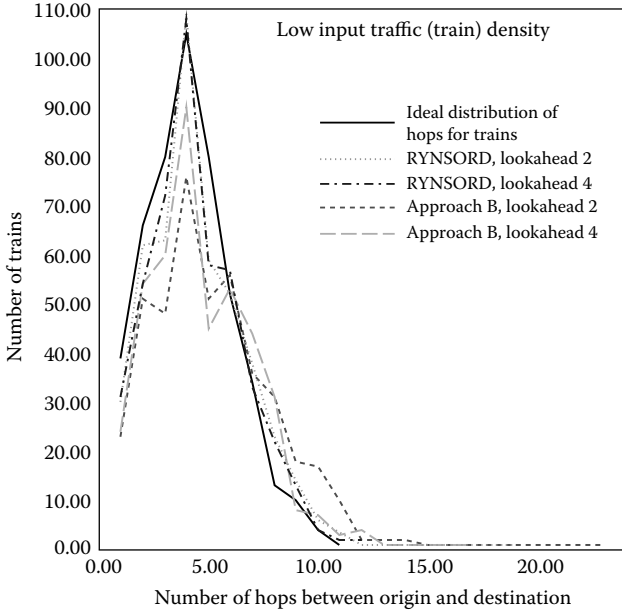
requests for the primary path. The stations, in sequential order, will try to reserve the requested track at the desired timesteps. If successful, the train uses the approved tracks of the primary path. Otherwise, if any of the tracks are busy, the reservation request is denied and immediately returned to the train at the originating station. Under these circumstances, the train then sends out hard reservation request on the secondary path. If this also fails, the train must wait a minimum of one timestep before initiating a request again to the primary path. This process continues until the train is able to acquire reservation and move forward. Conceivably, a train may have to wait at a station prior to succeeding in acquiring reservation approval. If the reservation request is successfully approved, the train moves along the  $N$  consecutive tracks.

The key performance measures include (1) the travel time of trains, averaged over all trains arriving at their destinations, (2) the percentage of trains reaching their destinations, (3) the distribution of the number of hops (tracks) utilized by the trains, (4) the average number of hops traveled by the trains as a function of the lookahead, (5) the travel time of individual trains as functions of the times of their assertion into the simulation, and (6) the track utilization. Furthermore, to understand the importance of distributing the overall computation and communication tasks among all entities, three additional performance measures are defined. They include (7) the distribution of computations performed by the trains, (8) the distribution of the numbers of reservations processed by the stations, and (9) the maximum communication rate of the interstation links.

Figure 5.9a presents a plot of the (actual travel time of a train minus its ideal travel time), averaged over all trains, as a function of the lookahead size. The ideal travel time of every train is used as a reference, and it refers to the travel time that a train would require if it was the only one in the entire system and could proceed to its destination along the shortest path, unhindered by any other train. Clearly, in the presence of other trains in the system, a specific train may not succeed in acquiring reservation for and traveling on the tracks along its shortest path. Figure 5.9a shows six graphs, corresponding to RYNSORD and approach B for each of the three densities. For low and medium input traffic densities, RYNSORD's performance consistently exceeds that of approach B. Figure 5.9a reveals that the average travel time of trains increases modestly with increasing lookahead size. For high traffic density, the relatively poor performance of RYNSORD compared to approach B is an aberration that may be explained by the graphs in Figure 5.9b. Figure 5.9b plots the percentage of trains reaching their destinations prior to the termination of simulation, for both RYNSORD and approach B and for all three densities. As the density increases, the consequent greater congestion is responsible for lowering the percentage of trains that are able to finish their journeys. Also, in every case, a greater percentage of trains reach their destinations under RYNSORD than approach B, implying the superiority of soft reservations. Furthermore, since significantly less number of trains reach their destinations in approach B under high traffic density, relative to RYNSORD, the corresponding travel time graph for approach B in Figure 5.9a fails to include trains that run to farther destinations and is therefore skewed.



**FIGURE 5.9** (a) Average travel time minus ideal travel time averaged over all trains as a function of the lookahead, (b) percentage of trains reaching their destinations within the maximum simulation time.

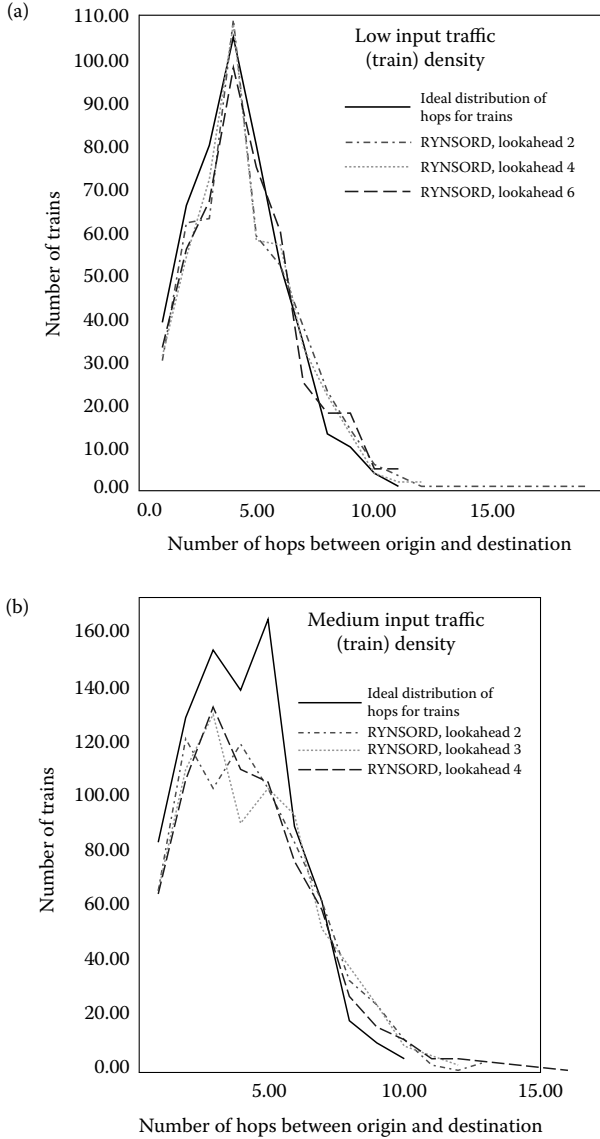


**FIGURE 5.10** Distribution of actual number of hops for trains in RYNSORD versus approach B versus ideal distribution.

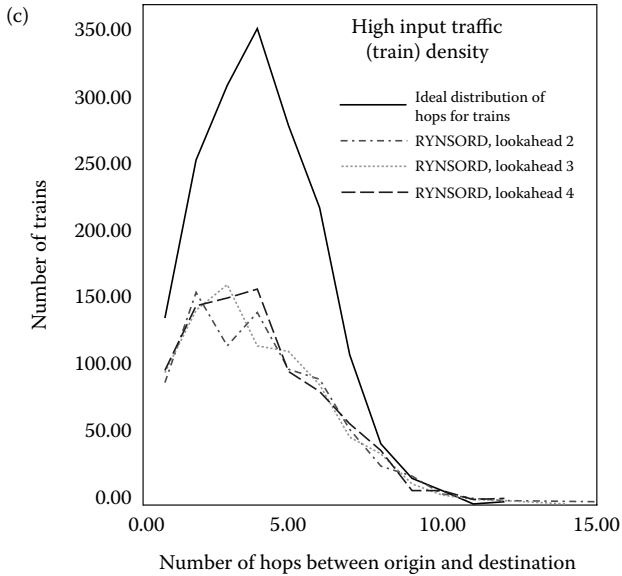
Figure 5.10 shows a plot of the hop (track) distribution of trains, that is, the number of tracks, ranging from 1 to 20, that are utilized by the trains to reach their destinations, corresponding to low input traffic density. Figure 5.10 shows five graphs, one corresponding to the ideal scenario, two relative to RYNSORD for lookahead values of 2 and 4, and two corresponding to approach B also for lookahead values of 2 and 4. The ideal scenario, described earlier, refers to the computation of the ideal paths that trains would take if every train was assumed to be the only one in the system. Under actual conditions, simulated on the test bed, it is highly probable that most trains will fail to acquire reservations for every track of their ideal paths since there will be demand for them from other competing trains. In sharp contrast, the graphs obtained from simulation show that the hop distribution closely follows the ideal scenario. That is, despite 484 trains competing for tracks, RYNSORD’s distributed, dynamic routing with soft reservation yields results that are close to ideal. The graphs are specially revealing for the following reason. There is a belief in the technical community that while distributed algorithms may yield faster results, in general, the quality of the distributed solution cannot approach that obtained from centralized algorithms. This belief is fueled by the fact that in distributed algorithms local agents execute the decision making but are allowed access to only a fraction of the system-wide data. The results from the rigorous RYNSORD simulation unquestionably refutes the generality of the belief. RYNSORD shows that under certain circumstances distributed algorithms may yield very high-quality solutions while generating them fast.



The authors are currently engaged in studying a new mathematical framework to extract distributed algorithms from centralized descriptions of problems. The graphs also reveal the superiority of RYNSORD's soft reservation over approach B's hard reservations.




**FIGURE 5.11** Distribution of actual number of hops for trains in RYNSORD versus ideal distribution under (a) low input train density, (b) medium input train density, and (c) high input train density.



**FIGURE 5.11** (continued)

The graphs in Figures 5.11a through 5.11c contrast the hop distribution of RYNSORD under different lookahead values relative to the ideal scenario for low, medium, and high traffic densities. The RYNSORD graphs in each of Figures 5.11a through 5.11c differ slightly from one another, implying that the impact of lookahead on the hop distribution is modest. Furthermore, with increasing traffic densities, the hop distributions increasingly deviate from the ideal scenario, implying that the increased competition for the tracks causes individual trains to select tracks other than those along their shortest paths from the origin to the destinations. In each of Figures 5.11a through 5.11c, the graphs corresponding to lookahead 2 reveals that a small, yet nontrivial, number of trains require excessive number of hops. This is due to double-backs, that is, where a train oscillates back and forth between two or more stations, while attempting to negotiate a suitable route to its destination. Despite the fact that this normally implies inefficiency of track usage, results from Figure 5.9a show that trains under lookahead 2, in general, reach their destinations faster. The occurrence of double-backs decreases substantially with higher lookahead values.

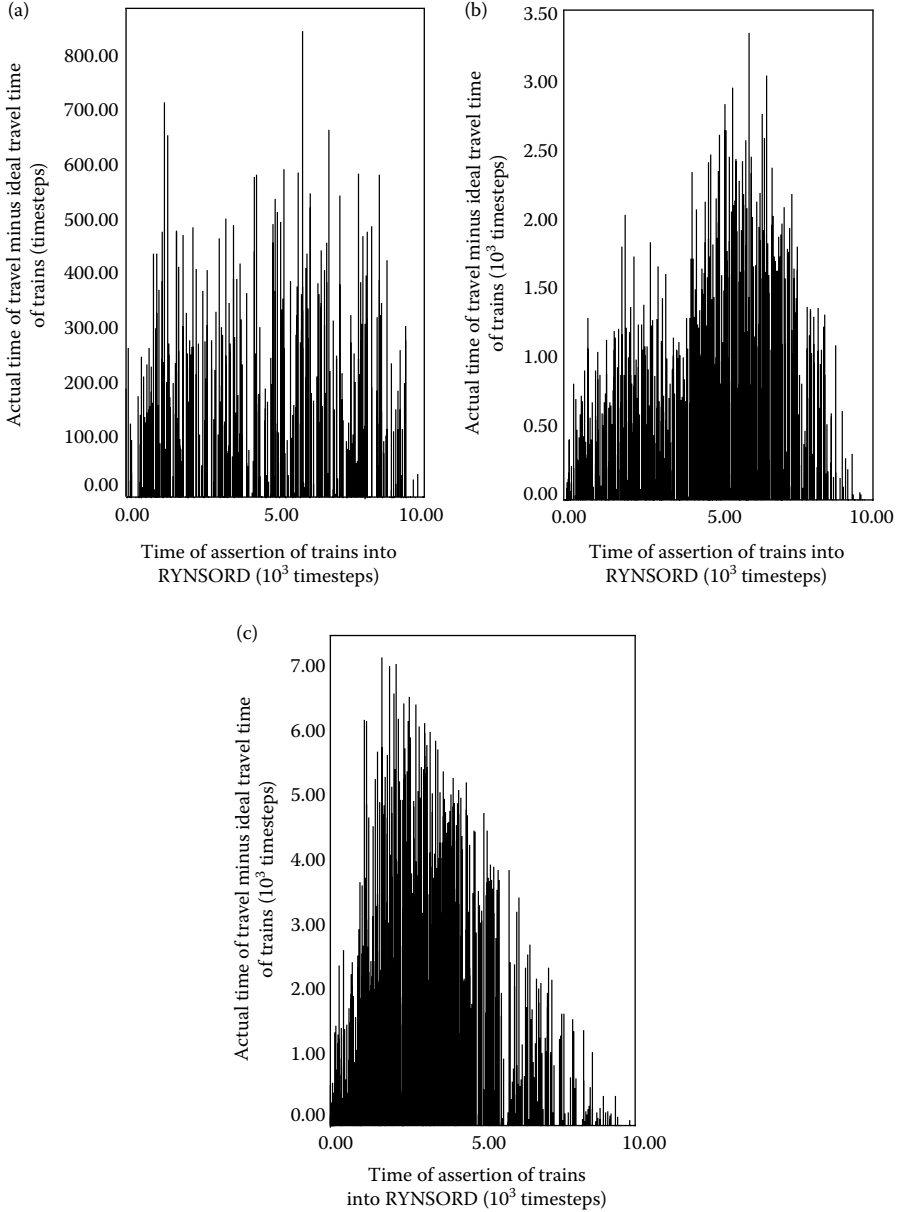
Figure 5.12 reports on the effort to study the impact of lookahead size on the average number of hops (tracks) for RYNSORD and approach B for each of the three input traffic density values. Once again, for every train, the ideal number of hops is computed and utilized as the standard against which actual number of hops utilized by the train is contrasted. The graphs for all three density values in RYNSORD appear to converge to a small value that is slightly lower than the corresponding value for approach B, once again demonstrating the superiority of soft reservations. Furthermore, for lower lookahead values in every case in Figure 5.12, the average



**FIGURE 5.12** Average hops minus ideal number of hops averaged over all trains as a function of lookahead.

number of hops relative to the ideal hop count is significantly higher. This corroborates the earlier finding that lower lookahead encourages frequent switching of tracks in the course of routing and trains traverse more tracks in the process.

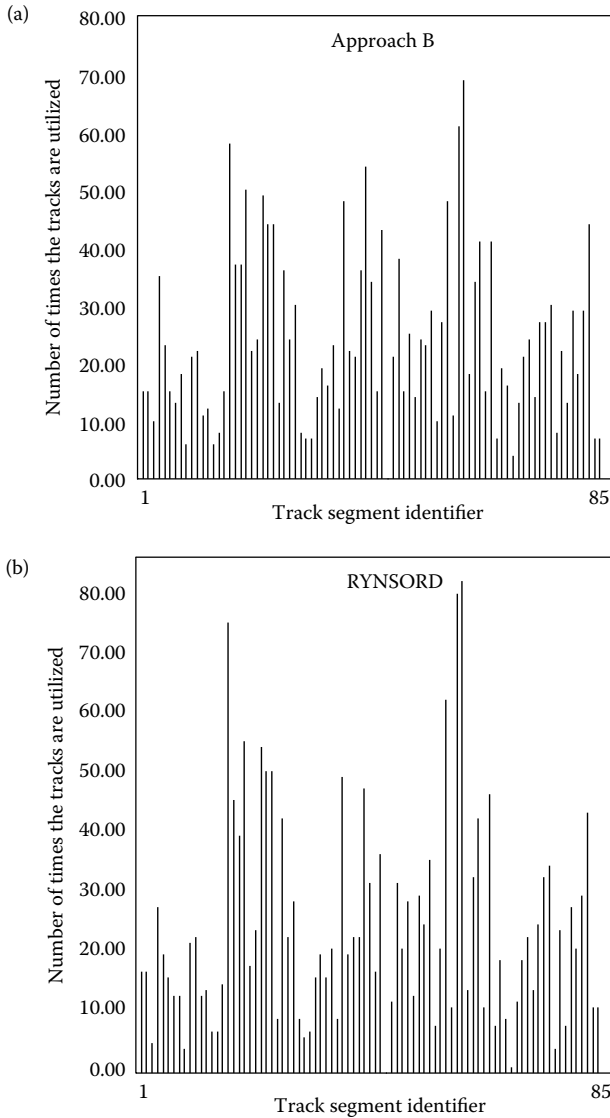
Figures 5.13a through 5.13c present the tuples [(actual travel time of a train minus its ideal travel time), time of assertion of the train into the system] for all trains that reach their destinations. Figures 5.13a through 5.13c correspond to low, medium, and high input traffic densities, respectively. In general, as more and more trains compete for tracks, trains will require longer time to reach their destinations. This is reflected by the increasing timestep scales along the Y-axes from Figure 5.13a to 5.13c. For low input traffic density, most trains reach their destinations regardless of their time of assertion into the system, and this is reflected by the relatively uniform distribution in Figure 5.13a. While the plot in Figure 5.13b exhibits modest cutoff at high values of assertion time, that is, 10,000 timesteps, that for Figure 5.13c is quite severe. This reflects the fact that under higher input traffic densities a train,  $T_a$ , that is asserted later into the system relative to another train,  $T_b$ , may require more time for travel than  $T_a$  and, under certain circumstances, it may not succeed in completing its



**FIGURE 5.13** Actual travel time of a train in RYNSORD minus ideal travel time for look-ahead 4 as a function of assertion time of the train, (a) low input traffic density, (b) medium input traffic density, and (c) high input traffic density.

journey within the maximum allowed simulation time. Clearly, to achieve a stable, continuous running system with minimal cutoff, one must select an appropriate input traffic density.

Figures 5.14a and 5.14b present track utilization results, that is, the cumulative number of times every track is utilized by trains, for approach B and RYNSORD.



**FIGURE 5.14** Track utilization for low input traffic density and lookahead 6, (a) approach B, (b) RYNSORD.

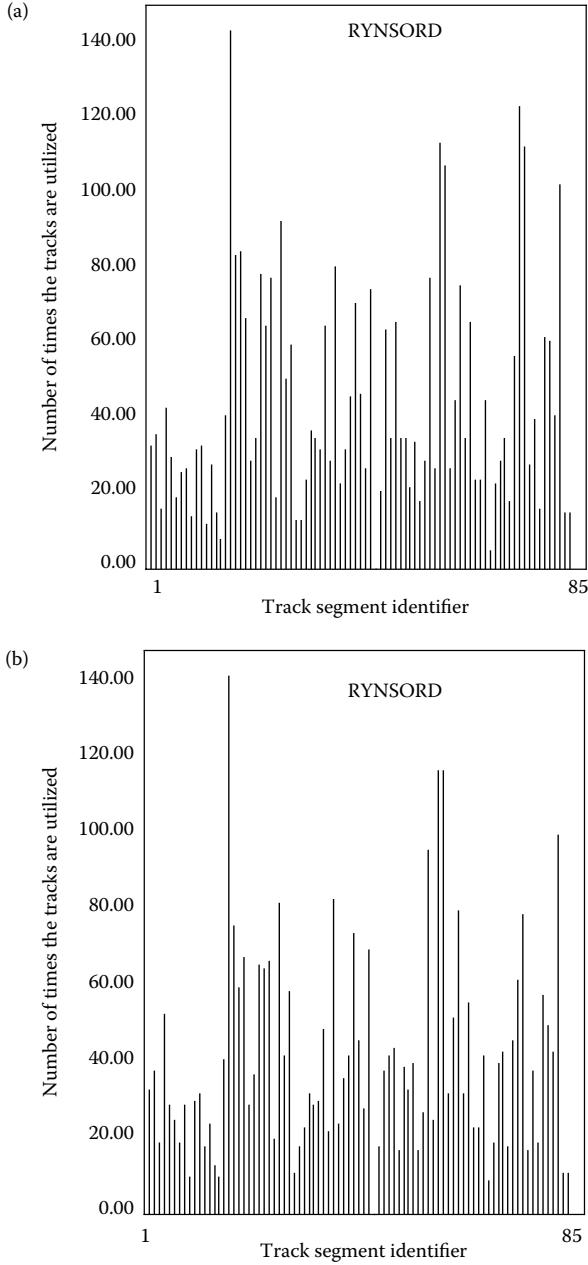
Track segments are identified by unique identifiers 1 through 84. While most of the tracks are utilized reasonably, reflecting efficient resource utilization, a few tracks exhibit high utilization, which merely reflects the stochastic destinations of trains and the choice of the “high traffic” stations. The track utilization plot in Figure 5.14b is observed to be, in general, higher than that in Figure 5.14a, implying the superiority of RYNSORD’s soft reservation over hard reservation in approach B.

Figures 5.15a and 5.15b present the track utilization for high input traffic density, and, clearly, it is significantly higher than that for low input traffic density in Figure 5.14b. However, the track utilization in RYNSORD is not significantly affected by the choice of the lookahead value.

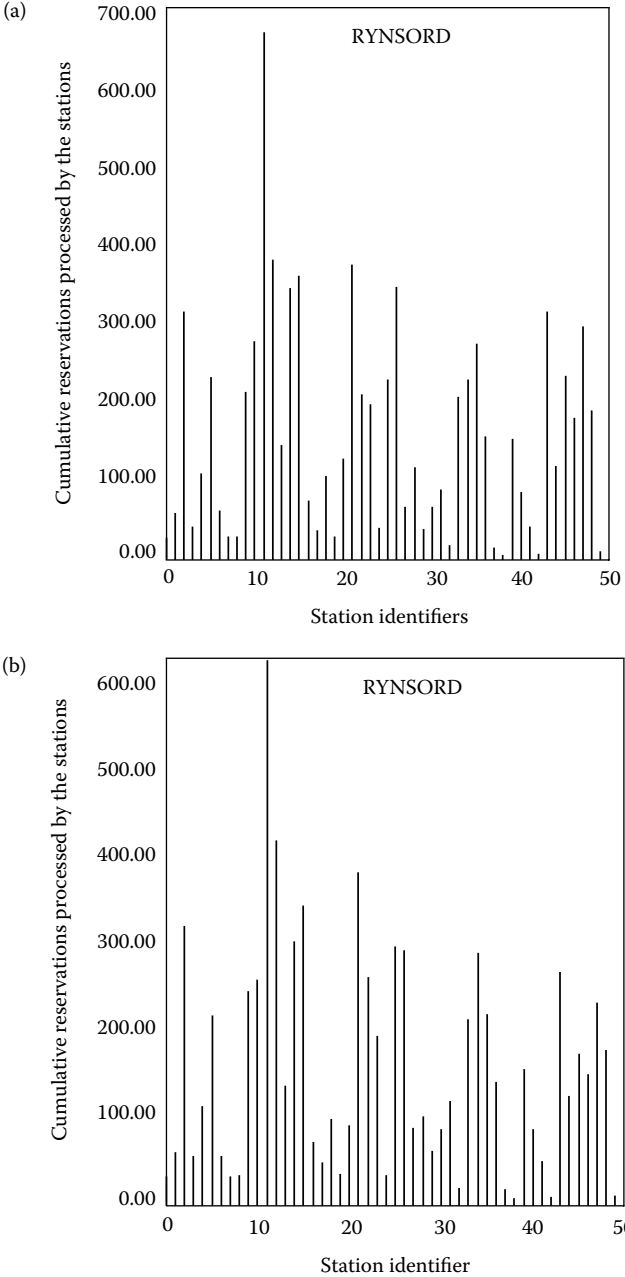
Figure 5.16 presents the distribution of a part of the overall computation task of routing the trains among the stations. A principal component of the overall computation task is reservations processing. While Figure 5.16a corresponds to lookahead 2, Figure 5.16b relates to lookahead 4. The computational load distribution among the stations is slightly higher for lookahead 2 than lookahead 4. Although the individual trains, under low lookahead, execute the Dijkstra’s [44] shortest path computations more frequently, they reserve fewer stations at any given time. Both Figures 5.16a and 5.16b underscore the achievement of the original goal of efficiently distributing the overall task among the stations nodes. The nonuniform distribution of the reservations processing is due to the stochastic destinations of trains, which, in turn, affects their routing.

Figure 5.17 presents the distribution of the remainder of the overall computation task, among the trains. A principal component of the overall task is the Dijkstra shortest path algorithm execution by the trains for computing the primary and secondary paths. While Figure 5.17a corresponds to lookahead 2, Figure 5.17b relates to lookahead 4. Both Figures 5.17a and 5.17b underscore RYNSORD’s original goal of efficiently distributing the overall task among the trains. In both Figures 5.17a and 5.17b, with the exception of a few trains, the computational burden is uniform among most trains, which underscores the achievement of equitable distribution of the overall task among all trains. The computational load in Figure 5.17a is significantly higher than that in Figure 5.17b since, under low lookahead, trains perform shortest path computations more frequently.

Tables 5.2 and 5.3 present data collected from the simulations to assist in the understanding of the impact of lookahead on key performance measures. The “average time” refers to the travel time of trains relative to the ideal travel times and is computed as equal to  $\{\text{sum overall trains of (actual travel time of a train – its ideal travel time)}\} \div \text{by the total number of trains}$ . While Table 5.2 presents data for low input traffic density, Table 5.3 corresponds to medium input traffic density. In Table 5.2, for high lookahead values, the average travel time of trains increase, and it correlates to the commensurate increase in the average waiting time. The latter, in turn, is due to the fact that to reserve more tracks for increasing  $N$ , trains must wait longer at the host stations where they initiate reservations, while engaged in communicating with more stations. However, the average number of hops decreases with increasing lookahead while the decrease in the number of double-backs is even more dramatic. The contrast between low and high lookaheads is more pronounced in Table 5.3.

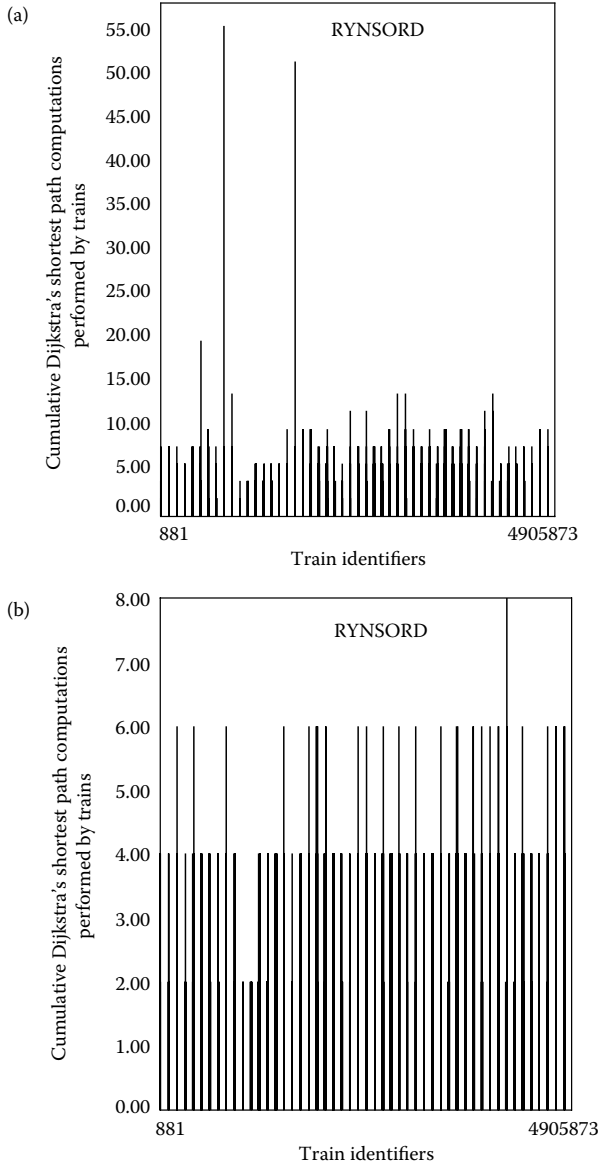


**FIGURE 5.15** Track utilization in RYNSORD for medium input traffic density for (a) lookahead 2, (b) lookahead 4.



**FIGURE 5.16** Distribution of reservations processed by stations in RYNSORD for medium input traffic density as a function of stations for (a) lookahead 2, (b) lookahead 4.





**FIGURE 5.17** Distribution of computation among the trains in RYNSORD for medium input traffic density for (a) lookahead 2, (b) lookahead 4.

For low lookahead, the average travel time of trains and average waiting time are significantly lower. However, the average miles traveled by trains, the average number of hops, and the link usage are higher. In addition, the number of double-backs is particularly high. Trains under low lookahead have a restricted view of

**TABLE 5.2**  
**Comparative Impact of Lookahead on Performance Parameters in RYNSORD under Low Input Traffic Density**

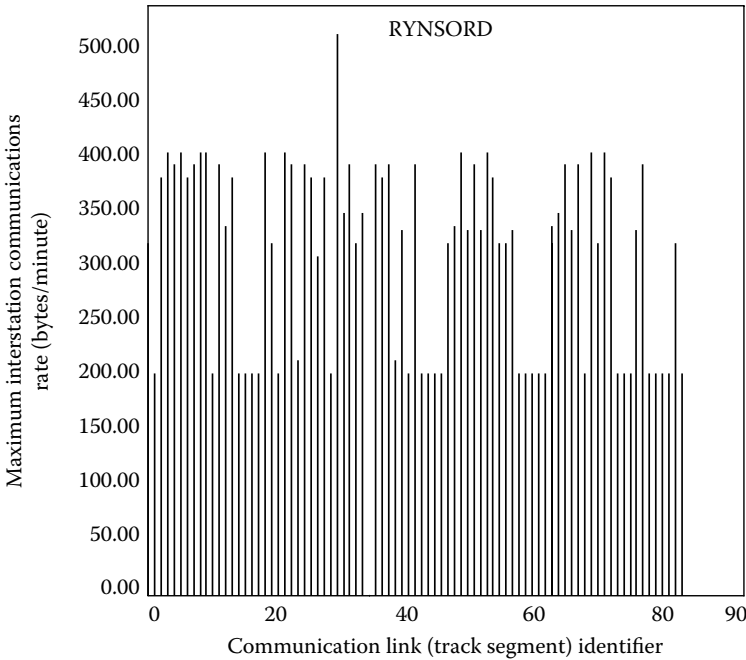
Lookahead Value	Link Usage (%)	Total No. of Hops	Total No. of Double-backs	Average No. of Hops per Train	Average Time	Average Miles per Train	Average Waiting Time (timesteps)
1	27	2313	321	5.039	122.11	620	113
2	27	2053	17	4.483	151.91	627	119
3	27	2026	5	4.443	164.24	619	137
4	27	2029	3	4.450	170.56	627	136
5	27	2046	2	4.458	182.01	620	151

**TABLE 5.3**  
**Comparative Impact of Lookahead on Performance Parameters in RYNSORD under Medium Input Traffic Density**

Lookahead Value	Link Usage (%)	Total No. of Hops	Total No. of Double-backs	Average No. of Hops per Train	Average Time	Average Miles per Train	Average Waiting Time (timesteps)
1	52	4656	1274	6.151	545.94	647	615
2	48	3428	144	4.565	688.95	619	647
3	46	3194	12	4.357	759.77	610	711
4	45	3186	9	4.341	769.62	604	728
5	45	3137	3	4.274	780.01	603	737

the system-wide congestion and are more likely to make poor long-term choices. However, they make routing decisions more frequently, and although this increases their computational burden, their decisions are up to date and superior as reconfirmed by the shorter average travel times. In contrast, trains under high lookahead are locked into tracks for longer periods of time and fail to take advantage of rapid dynamic changes in the system-wide track usage, as reflected by their longer average travel times. Nevertheless, their routing is more organized, requires less hops and distance traveled, and virtually eliminates double-backs. Thus, where shorter travel times are of paramount importance and the cost of track usage is negligible, low lookahead values are logical. Where the cost of using tracks is appreciable relative to the idle waiting of trains at stations, high lookahead value is recommended.

A principal objective of RYNSORD is to minimize the intra-network communications by distributing the overall computational task to the local entities. [Figure 5.18](#)



**FIGURE 5.18** Maximum communications rate (bytes/min) of interstation links.

presents a plot of the maximum communications rate for every interstation communication link. Given that the resolution of the simulation is 1 timestep or 1 minute of real operation, the resolution of the data presented here is also limited to 1 minute. A maximum of 500 bytes/min of data propagation is observed in Figure 5.18, which is easily realizable through commercial wireless modems rated at 9,600 or 19,200 baud. Thus, one of the principal objectives of RYNSORD is achieved. In contrast, a centralized algorithm would theoretically require a much higher communications rate, implying expensive interfaces.

*Limitations of RYNSORD*

One limitation of the current RYNSORD implementation is that it does not model abrupt track failures. Conceivably, track failures may cause severe local congestions, which may spread to other parts of the network. While RYNSORD allows trains to utilize congestion information to replan their routes, its performance in the event of track failures warrants further study, and is reported in Chapter 7.

---

# 6 DICAf: A Distributed, Scalable Architecture for IVHS

## 6.1 INTRODUCTION

According to the Intelligent Transportation Society of America (ITS America) [24], surface transportation in the United States is at a crossroads. While the nation's 4 million miles of paved roads are badly clogged with 200 million cars and congestion continues to increase, the conventional wisdom of building more roads will not work for both financial and environmental reasons. While congestion costs 100 billion dollars (Frank Kreith, pers. comm., May 16, 2000) annually in lost productivity, energy wastage, vehicle deterioration, road rage, and increased emissions from vehicle idling, traffic accidents in 1993 alone caused 40,000 deaths and 5 million injuries. In response to these problems, the U.S. Congress passed the Intermodal Surface Transportation Efficiency Act (ISTEA) of 1991, whose basic goal is to develop a national transportation system that is economically efficient, environmentally sound, and moves people and goods in an energy efficient manner. The U.S. Department of Transportation (DOT), led by the Federal Highway Administration, has launched the Intelligent Vehicle Highway System (IVHS) program to meet the demands of the ISTEA. IVHS does not aim to address the capacity problem. It aims to assist in steering drivers away from bottlenecks and in introducing and managing reasonable enforcement measures such as congestion pricing. Surface transportation-related problems are not unique to the United States. In fact, in countries with higher population densities such as Europe and Japan, the problem is more acute. The Programme for a European Traffic with Highest Efficiency and Unprecedented Safety (PROMETHEUS) [45] project in Europe and the Advanced Mobile Traffic Information and Communication System (AMTICS) program [46] in Japan closely parallel the IVHS program in the United States.

To virtually every driver today, the current interstate and state highway system is often a source of frustration, the principal reason being congestion. King [47] notes that driver navigational waste is equal to 6.4% of all distance and 12% of all time spent in travel by noncommercial motorists, amounting to millions of dollars. The Rhode Island Department of Transportation (RIDOT) [48,49] estimates that 60% of all vehicle-hours lost is due to accidents, stalled vehicles, and other road mishaps, which are dynamic and unpredictable. Similar findings are reported by other State DOT agencies. Peters, McGurrin, Shank, and Cheslow [50] estimate that the ITS infrastructure must improve vehicle handling capacity by 30% in order to keep congestion from growing beyond the current level. Other sources of congestion include routine maintenance, construction, and special events, which are mostly predictable.

It is generally and probably correctly believed that availability of accurate highway-related information may constitute an effective antidote to such problems.

Consider the following three typical scenarios: (1) A driver enters into a highway near exit 24, for example, and runs into a congestion within a mile. The congestion is severe but extends only up to exit 22, for example. Had the driver been aware of this problem before entering into the highway, he/she could have easily driven on a short back-road and entered the highway past the congestion. (2) A driver passes a gas exit and after driving 20 miles notices that the fuel level is low and that most gas exits are closed. Had the driver been aware of this problem, he/she could have easily filled the tank at the last gas exit. (3) A driver needs to reach home in Rhode Island from New Jersey on a Sunday. It is 1 PM in the afternoon and it has just started to snow. The forecast calls for heavy snow after 7 PM. The driver assumes that I-95, an important thoroughfare, will be kept cleared and that it would require approximately 5 hours to drive the 240 miles, yielding an estimated time of arrival of 6 PM. The driver enters interstate I-95, only to find out that the interstate is not being cleared of snow and that it is impossible even to pause in the breakdown lane to clear the icy rain from the windshield. The average speed hovers around 25 mph and the driver is caught in the heavy snow storm only to reach home after a painful 10 hours. In each of the above scenarios, clearly, if accurate information was made available, the drivers could have judiciously planned their trips and thereby avoided contributing to the congestion, while ensuring safety and economy.

In addition to substantial research reports on traffic management and traffic control, the recent literature reports a number of research efforts in the ITS. Haver and Tarnoff [51] report a new, efficient traffic management system that utilizes micro-processors and local area networks to achieve online signal optimization. Fenton and Mayhan [52] report their studies and findings relative to the development of theoretical control concepts and controllers for longitudinal and lateral control toward an automated highway system. Powell [53] presents a summary of current tools used in the optimization of assignments of drivers to motor carriers, dynamic fleet management, that is, pickup and delivery schedules, and vehicle routing. He observes that the traditional vehicle routing problem is a fundamentally hard, mathematical problem, and that there is little difference between the different techniques that are in use today. Batz [54] reports on the use of TRANSCOM, the real-time traffic information that is dispatched by over 14 transportation and traffic enforcement agencies in the New York/New Jersey metropolitan area and utilized by the trucking industry. When an incident occurs, TRANSCOM reports the location and time of incidence and an estimated time to clear through a 80-character message. This is intercepted by the participating motor carriers, which then analyzes the impact on their individual trucks and relays appropriate information on a timely basis. While the system is currently under evaluation, it is expected to facilitate increase in fleet productivity, driver's environment, and customer service. Roper and Endo [55] report on the Santa Monica Smart Corridor Project, whose primary objective is to create a better balance of flows among all roadway facilities. It is proposed to develop a centrally located urban freeway traffic control center, Central, that will collect traffic data from multiple

sources and disseminate them, in real time, to in-vehicle displays on 25 selected, en route vehicles. Once every minute, the in-vehicle processor receives messages containing link congestion data that are broadcast from the Central, and uses it to extract relevant congestion data that is pertinent to the location and heading of the vehicle. At the Central, a dedicated workstation tracks every one of the test vehicles to monitor vehicle routes and diversions made by drivers. This scheme is expected to lower the million vehicle-hours per year by 15%, the average freeway trip duration by 12%, and increase the average freeway speeds from 15–35 mph to 40–50 mph.

In its incident management plan [48,49], the RIDOT aims to inform the public of predictable and dynamic sources of congestion through traffic reports on radio stations. For accurate reports, RIDOT plans to consolidate information from visual air traffic patrols, video cameras, RIDOT ground vehicles, public safety patrols, emergency vehicles, and motorists through a standardized information exchange format and by utilizing a combination of computers, modems, and fax machines. A 24-hour toll-free telephone incident reporting mechanism and a free “SP” dedicated cellular phone line are also planned. In addition, RIDOT plans to improve the use of existing traffic loops and explore alternate detection schemes in high-accident-prone areas.

Kremer, Hubner, Hoff, Benz, and Schafer [56] present a short-range mobile radio network, referred to as mobile radio LAN’s, for IVHS and describe a simulator, MONET3, that allows evaluations of protocols to operate in networks with hundreds of stations. A key advantage is that most traffic data may be provided locally without the need for global communications. Sakagami, Aoyama, Kuboi, Shirota, and Akeyama [57] describe a methodology to determine vehicle position in multipath environments from the angle of arrival of waves received by multibeam antennas. While the accuracy is hindered by tall building, the approach may be highly beneficial to track lost vehicles. Hussain, Saadawi, and Ahmed [58] describe a mechanism to detect and monitor traffic through an experimental overhead infrared optical system. The system successfully detects and counts vehicles and is weather-resistant and cost-effective. Kim, Liu, Swarnam, and Urbanik [59] describe a area-wide traffic control system (ARTC), wherein traffic flow information is frequently exchanged between signal controllers to successfully address frequent occurrences of congestion. The system exhibits improved success over an optimized fixed time control and adequate level of fault tolerance. In a related discipline of automobile highway, Von Tomkewitsch presents ALI-SCOUT [60], a dynamic route guidance system with onboard computers. An automobile receives routing information from a centrally located traffic computer through infrared communications beacons that are strategically located at traffic lights. The central computer utilizes current traffic conditions to determine a route tree, that is, the best routes. The onboard computer receives the route tree and selects the appropriate route based on its destination. The report superbly discusses key issues relative to the use of infrared communications beacons and notes that the approach had been field tested for 700 vehicles. However, it does not provide any performance measures, and, while it is uncertain whether the approach will scale-up, the use of a central computer to generate the route tree is likely to inhibit the scalability of ALI-SCOUT. Denney and Chase [61] report that the use of distributed processing in

the San Antonio downtown traffic system has resulted in an open architecture that is responsive yet cost-effective and reliable. Kline and Fuchs [62] report that while the visibility of symbolic highway signs is significantly higher than those of same-sized text it may be greatly enhanced through the use of improved symbolic signs designed using an optical blur (i.e., low pass) approach in order to avoid higher spatial frequencies. Robertson and Bretherton [63] describe the SCOOT method of optimizing traffic signals in real time that adapts the signal timings automatically to new flow patterns. Bernard [34] proposes the ASTREE railway traffic management system, which maintains a distributed database of up-to-date, accurate, and comprehensive representation of route layout and train progress. However, the centralized decision making in ASTREE uses the information in the database to either automatically make decisions or assist human operators with decisions, relative to route settings and train control. The settings are then downloaded to the wayside equipment and locomotives.

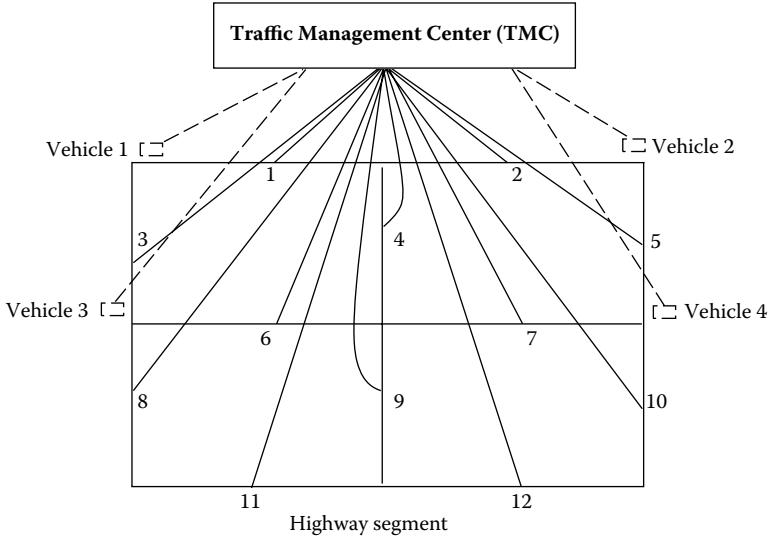
Researchers at Partners for Advanced Transit and Highways (PATH), at the University of California, Berkeley, have proposed an architecture for IVHS [26], wherein one or more automobiles are organized into discrete platoons that move through special lanes, similar to high occupancy vehicle (HOV) lanes, on existing freeways at very high speeds. When a vehicle enters into the network and announces its ultimate destination, the IVHS system assigns it a nominal route through the network. While the approach has been successfully tested [64], its limitations include the risks of entering and exiting the HOV lanes in the presence of other lower speed vehicles on the freeway and the fact that many drivers may resent the idea of being forced to travel at very high speeds. Shladover and colleagues [65] summarize their accomplishments relative to automating vehicle lateral (steering) and longitudinal (spacing and speed) control. von Aulock [66] reports that a feasibility analysis of automatic guidance system on German freeways conducted by Prof. Hiersche of the Technical University of Karlsruhe has concluded that while existing freeways and bridges are not built and cannot be modified economically for automatic vehicle guidance system, entering and leaving the system with vehicles zooming along 5 meters apart at speeds reaching 120 km/h is a highly likely source of accidents.

ITS, America, has recently published [24] the design proposals released by the four national architecture development teams led by Hughes Aircraft, Loral Federal Systems, Rockwell International, and Westinghouse Electric. The Hughes approach consists of a centralized, traffic management center (TMC) that detects and analyzes incidents and mitigates congestion by issuing real-time traffic information, routing parameters, and through controlling ramp metering and traffic signal timing. The TMC is aided by area processors that principally control the communication between the TMC and the beacons that interact with the vehicles. The Loral approach utilizes the concept of a fully integrated transportation system, allows for modular and flexible subsystems, and supports open standardized interfaces. Rockwell proposes a multilayered architecture and recommends interface standards at the application layer for each interface. It also proposes the development of a TMC but deliberately avoids specifying the design and configuration of the TMC. The

Westinghouse team proposes the use of the TMC and traffic control centers to provide centralized route guidance and other information to vehicles. A key concern with the proposed architectures is that they are yet to be supported by scientific experimentation, mathematical validation, or simulation. Under contract from the Federal Highway Administration (FHWA), Nynex Corporation [67] has developed a traffic management system, wherein a centralized TMC serves selected test vehicles that are equipped with specialized cellular phones. Initial data from the “operational tests” show that the length of an average cellular phone call necessary for a TMC to provide route guidance instructions to a single vehicle is approximately 5–10 minutes. The total number of vehicles in the NYNEX study is extremely small. While NYNEX estimates that it is too costly to gain widespread consumer acceptance, a more serious problem is that the length of the call will increase significantly as more and more vehicles demand interaction with the TMC. Studies by VanGrol and Bakker [68] in Germany corroborate NYNEX’s finding in that centralized TMCs that perform dynamic traffic management and short-term traffic forecasting are increasingly unable to keep up with the demand. Iqbal, Konheim, and Ketcham [69] note that the accuracy of the projections of recurring and nonrecurring congestions along corridors is limited by the static nature and highly variable quality of existing data. Jing, Huang, and Rundensteiner [70] recognize the difficulty of simultaneously computing a large number of paths for a huge transportation network and in storing the large number of precomputed paths in the computer memory. They propose a hierarchical encoding of the partial paths that offers improved performance and space efficiency. Talib, Love, Gealow, Hall, Masaki, and Sodini [71] propose incorporating a special, high-density pixel parallel processor chip onto a desktop computer to achieve fast low-level image processing in ITS systems. Ziliaskopoulos, Kotzinos, and Mahmassani [72] present techniques to execute shortest path algorithms, fast, on CRAY supercomputers. Rouphail, Ranjithan, El Dessouki, Smith, and Brill [73] propose the development of a decision support system for pre-trip route planning that generates maximally different routes for a network that is characterized by time-dependent link travel times. Centralized TMCs are complex and expensive to build and maintain. For instance, the Minneapolis TMC [74] contains 48 17-inch monitors, controls 354 ramp meters, receives data from 142 CCTV cameras that are located along the highways and connected through fiber optics, and managed by 37 personnel. Upchurch, Powell, and Pretorius [75] describe the deployment of a closed-circuit television camera network along 256 km of arterial corridors in Phoenix, Arizona, at a cost of \$42 million.

Chang, Junchaya, and Santiago [76] describe a traffic simulator implemented on a connection machine CM-2 and note that its performance is promising. Junchaya and Chang [77] state that their simulator has the inherent path processing capability to represent driver’s route-choice behavior. They report being able to simulate 32,000 vehicles for 30 minutes at one-second intervals in 3.5 minutes with 16,384 processors. Given that CM-2 is a SIMD (single-instruction multiple-data) architecture, that is, every processor executes the same instruction in lock step, and the processors are extremely simple; the simulator is incapable of modeling the complex, concurrent,





**FIGURE 6.1** IVHS architecture utilizing centralized traffic management center.

autonomous, and unique behavior of the individual vehicles. Also, SIMD machines are in essence synchronous machines and, therefore, they are not naturally suited to model the real-world, asynchronous, traffic system.

To illustrate the traditional efforts, consider a highway system, shown in Figure 6.1, that consists of 12 highway segments labeled 1 through 12, and the centralized TMC. The TMC is connected to each of the 12 segments through permanent links, shown through solid lines, that carry status information to the TMC as well as ramp metering commands back to the highway segments. The frequency of information exchange is governed by the flow of traffic, and it defines the accuracy of the TMC’s knowledge of the system state at any given time. When a vehicle, labeled vehicle-1 in Figure 6.1, enters the system, it establishes a temporary communication link with the TMC, shown through dotted line, and informs the latter of its final destination. The TMC then determines the route, taking into account the number of vehicles in the system, their destinations, and its goal of balancing the use of resources against the shortest travel times of the vehicles, and imposes it on vehicle-1.

A significant limitation with all of the above efforts is explained as follows: it must be recognized that it is neither logical nor feasible for a single, centralized TMC to continually broadcast every piece of highway data that any of the thousands of vehicles on the road may desire to know. Moreover, a single centralized TMC possibly cannot serve a wide geographical area effectively, given the limited power of radio transmitters, uneven terrain, and other factors. Furthermore, it is a well-known fact that a single, centralized unit cannot maintain the most precise information on the status of every highway segment at all times. Accuracy and precision is best achieved through a number of relatively autonomous and communicating local units. If the U.S.

DOT bases the IVHS architecture on centralized TMC, this chapter hypothesizes that the increase in the number of vehicles, the associated increase in congestion, and the increased demand for sophisticated traveler services and other highway-related information by drivers in the future will possibly require the total redesign of the IVHS system. The reason lies in the fundamental limitation of a single-processor computer. A TMC is a serial computing entity, that is, it executes its subtasks one at a time, and no matter how sophisticated and powerful it may be, its performance is bound to deteriorate as the number of vehicles interacting with it increases.

An additional limitation is that current efforts call for the TMC to exclusively divert and control the flow of traffic, based on its knowledge of congestion of all relevant highways. While this may be beneficial under certain scenarios, at other times, this may evoke resentment from independent-minded drivers. There are additional problems. First, as explained earlier, it is certain that a TMC will not have accurate and up-to-date information on all highway segments at all times, particularly when the numbers of segments and vehicles are large. Second, the TMC dictated alternate route to a vehicle may not be ideal since the TMC cannot consider reasons that are unique to each and every driver. For example, assume that a TMC dictates a driver, traveling through Rhode Island on I-95 and currently near exit 5, to take I-495 so as to divert traffic from an accident around exit 14. On the contrary, the driver who is tired from driving all night would have opted to travel forward up to exit 12, enjoyed breakfast at a restaurant for an hour, and then driven forward, had he/she been advised of the problem rather than dictated a rerouting. By that time, the accident would have been cleared. In this nation of independent-minded individuals, technology that offers choices is preferred to those that dictate one-size-fits-all type solutions. Kawashima [46] observes that a unified AMTICS and RACS in Japan may realize route guidance through onboard computers and not by central computers but notes that this is merely a conceptual model. Dailey, Haselkorn, and Meyers [69] correctly observe that a key element in ITS is the distribution of dynamic data in real time to a large but authorized group of users. They propose the use of an asynchronous, distributed, client/server architecture that relies on the creation of autonomous, reusable pieces of hardware. Hall [78] argues that advanced traveler information systems must aim at utilizing alternate routes, where possible, to steer traffic away from disequilibrium behavior and to provide to the user confidence and comfort in the system.

Recently, in a number of U.S. cities, notably Seattle, Houston, and Los Angeles, the congestion information on different interstate and highway segments are collected periodically and displayed on the Internet [79]. A total of five categories (George H. Way Jr., pers. comm., February 1992) are utilized to overlay the congestion information on a map namely, wide-open, heavy, moderate, stop-and-go, and no data available. While the practical difficulty of automatically acquiring this information onto every vehicle needs to be addressed, the approach supports a rudimentary requirement of DICAF, the central theme in this chapter. A serious issue, however, is the latency of this information, that is, the difference between the times that the information is generated and utilized, and its impact on the timeliness and accuracy of this highly dynamic information.

This chapter recognizes these problems and proposes the use of a distributed strategy, DICAf. DICAf utilizes the same basic principles, namely that of asynchronous, distributed algorithms, as presented in Chapters 4 and 5. The principles consist of utilizing data locally to compute decisions, wherever possible, and in propagating changes, that is, new information, to other entities in the system on a need-to-know basis. Clark and Daigle [80] review the vital importance of computer simulation in traffic engineering and stress its critical role in the development and evaluation of new ideas, algorithms, and traffic control systems. The remainder of this chapter is organized as follows. Section 6.2 details the DICAf algorithm while Section 6.3 describes the modeling of DICAf on an accurate, realistic, parallel processing test bed. Section 6.4 presents the details of implementing DICAf on the test bed. Section 6.5 reports on the simulation of representative traffic networks under stochastic and realistic input traffic and also presents a detailed performance analysis.

## 6.2 DICAf: A NOVEL, DISTRIBUTED, AND SCALABLE APPROACH TO IVHS

The vehicle routing problem is perhaps one of the richest problems in transportation [53] both because of its wide applicability and its fundamental complexity from a mathematical point of view. The Bodin et al. [81] review contains over 700 references over the past four decades, a testimony to the richness of the problem. As indicated earlier, it must be recognized that it is neither logical nor feasible for a single, centralized TMC to continually broadcast every piece of highway data that any of the thousands of vehicles on the road may desire to know. Moreover, a single centralized TMC possibly cannot serve a wide geographical area effectively given the limited power of radio transmitters, uneven terrain, and other factors. Furthermore, it is a well-known fact that a single, centralized unit cannot maintain the most precise information on the status of every highway segment at all times. Accuracy and precision is best achieved through a number of relatively autonomous and communicating local units.

DICAf recognizes these problems and proposes the use of a distributed architecture wherein the overall task of data collection, processing, dissemination of information, and decision making is distributed among all of the components of the IVHS system. The fundamental philosophy is to intelligently distribute decision-making tasks among the entities to maximize local computations, minimize communications, and achieve robustness and high throughput. A direct consequence of this philosophy is scalability, that is, where the underlying system will continue to function and deliver relatively undiminished performance as the system grows in size with an increasing number of vehicles and highway segments. The intent of the architecture is to influence every driver's routing decision by providing accurate and adequate highway data, timely, to help him/her plan alternatives, rather than dictate routes that inevitably lead to driver resentment and rejection of the system. Allen, Ziedman, Rosenthal, Stein, Torres, and Halati [82] report that, in simulation studies, the navigational system characteristics have significant effect on driver route diversion behavior, with better systems allowing more anticipation of traffic congestion.

Laboratory simulation studies with human drivers also indicate that the number of miles driven decreases when travelers are provided with better and more information.

To achieve the goals, this research effort requires (1) the understanding and analysis of the basic requirements of the entities—vehicles and road segments—constituting the traffic system, (2) identification of the essential information to be communicated between the entities, and (3) the determination of a generic model for all entities. The generic model must be capable of making independent decisions based on input information from a limited number of appropriate entities, yet each decision must be “consistent” with the others and cooperatively conform to the global goals of efficiency and safety. The proposed approach will carefully determine the basic rules of decision making and communication between the entities to address all possible scenarios. By definition, asynchronous, distributed algorithms offer, theoretically, the highest benefit from concurrent processing since there is no unnecessary interference nor synchronization. It allows the use of maximal distributed intelligence from the different entities. During the operation of an actual traffic system, the outcomes at different instances of time are functions of many parameters, such as the change of intent of a driver of a vehicle, vehicle malfunction, deterioration of road condition, accidents, and so on and cannot be predicted *a priori*. Since no single entity may possess accurate and complete knowledge of the entire traffic system at all instants of time, the asynchronous approach permits each entity to proceed as fast as it possibly can, without jeopardizing any aspect of the overall goal, namely efficiency of utilization of the resources and safety. The asynchronous approach also recognizes the intrinsic unique capabilities of every entity, if any, and permits their best utilization.

The current highway system is constituted by two components—highway segments and vehicles. A highway segment is simply a part of an existing highway, perhaps between two consecutive exits. The pavements/roadways analyzed in this chapter are freeways, without any loss in generality. While it is within a highway segment, a vehicle may not leave it to travel on an alternate route. Thus, when a highway segment, say between exits 2 and 3, is severely congested, a vehicle on the preceding segment may take exit 2 and travel on an alternate path. If it fails to take exit 2, the vehicle may not leave the segment until exit 3. The DICAF architecture introduces a third component—highway infrastructure, which is key to achieving its objectives. The infrastructure is organized into a number of constituent, Distributed Traffic Management Centers (DTMCs), each of which is responsible for the collection and dissemination of information within a well-defined locality. Every vehicle is assumed to be autonomous in that it is capable of requesting necessary information from the DTMCs, which it then utilizes to synthesize decisions. The information may assume many forms. Consider a vehicle, in transit, that requires very specific information. It communicates its request to the DTMC1, whose jurisdiction includes the current position of the vehicle. Such requests may either include local information such as the congestion information of a particular highway segment or nonlocal information such as the weather and driving condition of a highway segment significantly far away from its current position, or locations and business hours of banks, post offices, hospitals, and restaurants along the highway, further up from its current

position. Corresponding to a request for local information, DTMC1 immediately propagates the data to the requesting vehicle. When the information requested is outside its jurisdiction, DTMC1 may retrieve the data from the appropriate DTMC through the network and propagate it to the requesting vehicle. This chapter focuses only on the most fundamental parameter that is essential for route guidance, namely congestion information. The issue of retrieving information on the locations and business hours of banks, post offices, hospitals, and restaurants along the highway is beyond the scope of this chapter.

The exact location and number of the DTMCs in the DICAF system will be a function of (1) the number of vehicles requesting service, (2) the average electronic contact time between a vehicle and a DTMC, that is, the time needed to propagate the information, (3) the desired level of service, and (4) the range of the communication mechanism, that is, wireless or infrared beacon, between a vehicle and a DTMC. The principal choices in wireless communications mechanisms include the normal cellular telephones, cellular digital packet data system (CDPD), 220 MHz radio transmission [83] that has been set aside by the Federal Communications Commission (FCC) for FHWA's IVHS, and others. Kamali [84] presents a comprehensive review of a number of wireless communication technologies for the ITS. Sodeikat [85] reports that the short-range roadside infrared beacon, successfully demonstrated in the LISB field trials in Germany, supports up to 500 bits/s. The design of the inter-DTMC network will also be a function of the traffic volume and desired level of service. While a high bandwidth is likely to be necessary in an urban situation, a medium-bandwidth network will probably suffice for a rural community.

A driver's choice of the routing may be based on the weather, urgency, road condition, fatigue, the condition of the vehicle, and other objective and subjective issues. Although DICAF grants full freedom to every driver, for an objective evaluation of the DICAF algorithm, this approach assumes that every driver's basic objective is to reach the ultimate destination in the shortest possible time. There may be exceptions, especially for sight-seeing trips and other reasons. The driver's objective is not necessarily to utilize the shortest distance path since one or more segments of the path may experience greater congestion level. To assist the driver, this chapter records the following observations regarding the issue of congestion. As a first approximation, the average speed of vehicles on a given highway segment is a good measure since it reflects the throughput through the highway. However, the average speed does not capture the total number of vehicles on the highway segment, which, intuitively, must bear an impact on the congestion. Although the average inter-vehicle distance for a highway segment appears to reflect the total number of vehicles on the segment, by itself, it is neither a good measure of throughput nor congestion. It fails to differentiate between two scenarios where all vehicles are traveling at 50 mph and 100 mph respectively, while still maintaining the same average distance between the vehicles. For similar reasons, the total number of vehicles on a highway segment at any time instant is also not a good indicator of congestion. This chapter proposes a new definition of congestion measure (C.M.) of a highway segment, one that reflects the combined influence of the total number of vehicles on a highway segment and their

average speeds:

$$\text{C.M.} = \frac{\sum \text{Speed of vehicles}}{\sum \text{Number of vehicles}} \times \left(1 - \frac{X}{Q}\right) \quad (6.1)$$

In Equation 6.1,  $X$  represents the total number of vehicles on the highway segment at any instant. The maximum number of vehicles allowed on the highway segment,  $Q$ , is  $0.5 \times N \times \frac{D}{L}$ , where  $N$ ,  $D$ , and  $L$  refer to the number of lanes, the length of the highway segment, and the average length of a vehicle.  $D$  and  $L$  are generalized values. In the traffic engineering literature, two terms, “headway” and “gap,” measured and expressed in terms of time or distance, have been in use. While headway refers to the difference between the nose of the leading vehicle and the nose of the trailing vehicle, the gap is the difference between the rear of the leading vehicle and the nose of the trailing vehicle. For some analysis, especially under severe congestion as found on surface streets during rush hours and where the lead vehicle is a truck, these measures have been found useful. For this purpose of analyses in this book,  $D$  and  $L$  subsume headway and gap. Presumably, the highway segment can physically hold at most  $N \times \frac{D}{L}$  vehicles at any instant, all of which are traveling at the same speed. As an engineering approximation,  $Q$ , the maximum number of vehicles allowed on a highway segment, anytime, is assumed to be  $0.5 \times N \times \frac{D}{L}$ . The number of lanes for all segments, in this chapter, is assumed to be unity without any loss in generality. There is a general rule of thumb, usually found in the state driver manuals, of 1 car length separation between adjacent vehicles for every 10 mph speed. The reason this assumption is not reflected in the definition of C.M. above is that the maximum number of cars must reflect the physical maximum that a highway can hold at any time. This physical maximum is clearly defined by the length of the highway segment and the car lengths, and is not dependent on the vehicle speeds. The proposed definition in Equation 6.1 is more general in that the “maximum number of cars” encompasses even severe congestion scenarios. It also permits a much higher vehicle density than the rule of thumb would allow and this, in turn, enables the DICAF simulations to be driven hard to analyze its behavior under extreme conditions. The rule of thumb, as the term suggests, is only a rule of thumb. It neither has any formal basis nor is it strictly obeyed. It is also violated under severe congestion scenarios.

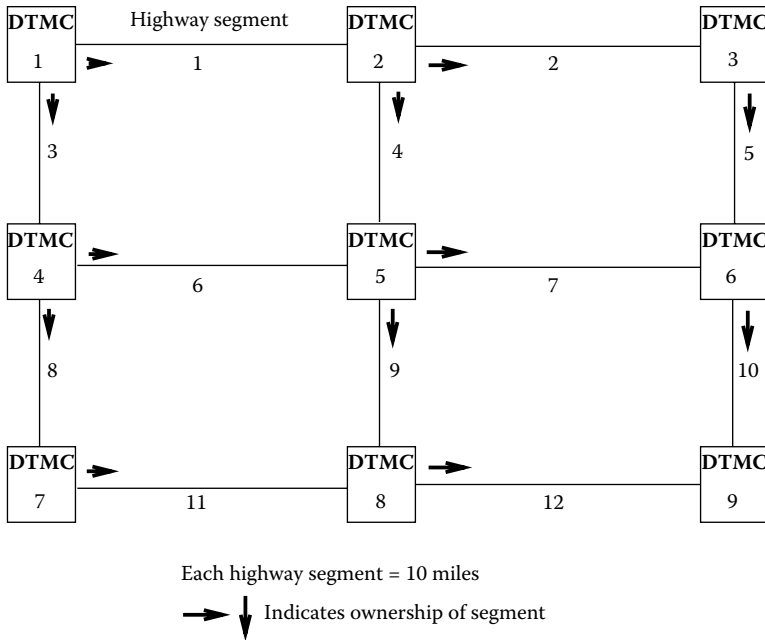
Additional rationale for the choice of the definition in Equation 6.1 is as follows: under normal circumstances, the speeds of the individual vehicles are different, as evident by the well-known fact that the lanes toward the left are designated as higher speed lanes while the right lane is meant for slower vehicles. In addition, the “passing” of lower speed vehicles by higher speed vehicles is a common occurrence in any highway. Now, when the number of vehicles in a segment is sparse, higher speed vehicles will be able to maintain their speeds even if a few vehicles are traveling at lower speeds for any number of reasons. Thus, the average speed is likely to be high. The situation changes dramatically when the number of vehicles increases and the inter-vehicle distance decreases. The higher speed vehicles find it increasingly difficult and unsafe to maneuver the slower speed vehicles to maintain their higher speeds. Consequently, they slow down, thereby bringing down the average speed

for that highway segment. Thus, the impact of increasing number of vehicles on a highway segment is that the congestion becomes severe and the C.M. value decreases. Where the average speed goes down to zero as is likely to be the case in the event of a severe accident, the C.M. value drops to 0. Also, when  $X$  equals  $Q$ , the C.M. value is 0. Since vehicles are assumed to travel at differing speeds, this condition implies that all vehicles have come to a stand still. This chapter assumes that the speeds of the vehicles follow a normal distribution curve with the “ $\mu$ ” and “ $\sigma$ ” values defined by the specific highway segment, the time-of-day, and so on. It is also pointed out that when the absolute value of C.M. is high, the level of congestion is low while a high level of congestion is implied by low values of C.M. While the highest C.M. value is defined by the maximum permitted speed on a highway segment, the lowest value is 0. Clearly, the definition of C.M. in Equation 6.1 will not apply in scenarios where all vehicles are traveling at the exact same speed. One could have used the well-known Poisson distribution. However, it is noted that while the Poisson distribution lends itself to easy mathematical manipulation, the DICA simulation is greatly facilitated by the readily generated tables from normal distribution. It is also pointed out that the simulations assume typical scenarios, that is, vehicles traveling year round as opposed to the time-of-day and day-of-year variation. Thus, the continuous normal distribution is utilized in this study without any loss in generality. Furthermore Kreyszig [86] notes that, for large  $N$ , that is, the number of independent performances of an experiment, a binomial distribution may be approximated by a normal distribution, while it is well known that the Poisson distribution may be derived as a limiting case of the binomial distribution.

As stated earlier in the chapter, most of the traditional route guidance approaches adopt a binary admittance policy, that is, they allow entry or refuse admittance, depending on the level congestion. It is well known that binary policies [87] generally lead to abrupt decisions. In contrast, in this chapter, the C.M. value evolves gradually as a function of the average speed and the number of vehicles on the segment. C.M. is a continuous function [86] in that, for any given segment, it is defined for and may assume any value in the range—(0 mph, maximum permitted speed on that segment). No vehicle is refused admittance into a highway segment, regardless of the level of congestion unless of course  $X$  is equal to  $Q$ . It is expected that vehicles that are not already within the segment in question will, in general, avoid segments with lower values of C.M. in favor of those whose C.M. values are higher. Thus, the overall routing decisions of the vehicles is expected to be less abrupt and the distribution of the vehicles overall segments of the highway system is likely to be more gradual and uniform.

Figure 6.2 presents the DICA architecture of a highway system that consists of 12 highway segments and 9 intersection points through which traffic is introduced into the system. DICA organizes the highway system into nine regions, each of which is controlled exclusively by its respective DTMC. Although the DTMCs are located precisely at intersection points in Figure 6.2, in truth, they may be located anywhere in the vicinity of the segments that they control. Thus, DTMC1 controls segments 1 and 3, while DTMC2 controls segments 2 and 4, and DTMC8 controls





**FIGURE 6.2** The DICAF architecture.

segment 12. DTMC9 does not control any highway segment. The DTMCs are connected through a wide-area network where the links parallel the highway segments. As soon as a vehicle enters the system, it communicates with the local DTMC to obtain the C.M. of other relevant segments. It then determines its own route based on the goal of reaching its destination quickly. The vehicle recomputes its route when it reaches the subsequent DTMC, and this process is repeated at every DTMC that it encounters, until it arrives at its destination. The C.M. value for a highway segment is most accurate within the DTMC that controls it and the accuracy decreases progressively, due to data latency, as one encounters other DTMCs that are further away. Thus, the vehicle progressively accesses accurate C.M. of the segments during its travel toward its destination. As a result, the vehicle continually refines its routing and achieves high efficiency.

Every vehicle contains a complete static topology of the highway system, that is, the number, length, and connectivity of the segments. However, vehicles lack knowledge of the C.M. values of the segments since they are dynamic. This chapter assumes that there is no permanent change in the static topology. The issue of road segment failure due to construction and incidents is beyond the scope of this chapter. When a vehicle is within the jurisdiction of a DTMC, it downloads from it the C.M. measures for the relevant segments. Then, it executes a modified Dijkstra shortest path algorithm [88] where the objective is to select a route that minimizes the estimated travel time (ETT) from the current position to the eventual destination. The ETT for



a segment relative to the vehicle in question, is obtained by,

$$ETT = \text{length of segment} \div \text{minimum (C.M., desired speed of the vehicle)} \quad (6.2)$$

In Equation 6.2, ETT is an estimate for the following reason: the C.M. measure for a subsequent segment is dynamic, and it may assume a value different from that used at the instance of computing ETT when the vehicle actually travels on that segment, if the vehicle does end up traveling on that segment. In addition, since C.M. is only an indicator of the congestion level, this chapter assumes the following: if the C.M. value is lower than a vehicle's desired speed when the vehicle is about to commence traveling on a segment, the vehicle is allowed to travel at a maximum speed of either 125% of the C.M. value or its desired speed. If the C.M. value is higher than the vehicle's desired speed, then, clearly, the vehicle is permitted to travel on the segment at its desired speed.

In turn, each DTMC computes the C.M. measures for the segments that it controls and propagates the values to other DTMCs using the flooding algorithm [88]. For efficiency, when a C.M. value for a highway segment differs from the previously propagated value by more than a predetermined fraction, only then is that value propagated. In addition, when a DTMC receives the C.M. values of other segments from other DTMCs, it updates its local record of the DICA system and propagates them in accordance with the principles in the flooding algorithm [88]. The appropriate information from its local record is propagated to a requesting vehicle within its jurisdiction. The DTMC utilizes its knowledge of the number of vehicles and their respective speeds on a segment at any given time instance to compute the C.M. Again, for efficiency, the computation is triggered when either one or more vehicles enter the segment or leave the segment. When a vehicle enters the jurisdiction of a DTMC and requests information, the DTMC registers the entry of the vehicle for the purpose of triggering the computation of C.M. When a vehicle leaves the jurisdiction of a DTMC1 and enters that of DTMC2, the latter will notify DTMC1 as soon as it registers the entry of the vehicle. Then, DTMC1 again retriggers the computation of the C.M.

The functions of each DTMC and vehicle is expressed, in pseudocode, in [Figures 6.3](#) and [6.4](#) respectively.

DICA is scalable, that is, as the system evolves and the number of vehicles and DTMCs increases, DICA continues to function, and its performance is expected to remain relatively undiminished. With an increase in the system size, the predominant task, namely that of computing the routing for the vehicles increases, which, in turn, is equitably shared by the proportionately higher number of computing entities. In contrast, under similar conditions, the sequentially executing central computer of the traditional TMC will be quickly overwhelmed by a significant increase in the computational burden. DICA is economical, that is, unlike the need for an expensive, central supercomputer at the TMC, each DTMC may be designed with a relatively inexpensive state-of-the-art microprocessor. The computational burden on a computing element of a DTMC is significantly lower. The maintenance cost of a DTMC computer, in addition, is also significantly lower. The DTMCs in DICA may be

```

While (simulation is not complete) {
  Check for flooding messages from neighbors
  If (a message with higher sequence number) {
    Update C.M. of segments for other DTMCs stored locally
    Propagate the message to other neighbors
  }
  Check for vehicles entering any highway segment it owns
  If (a vehicle has entered) {
    Communicate with the vehicle, download C.M. measures
    Upload vehicle's actual speed
    Determine when vehicle will exit segment based on its actual speed
    Add vehicle to the wait-to-leave list
    Update C.M. of the affected segment
    If (change in C.M. exceeds a specified threshold) {
      Send flooding message to other neighbors
    }
  }
  Check the wait-to-leave list
  If (a vehicle exits the segment) {
    Update the C.M. of the affected segment
    If (the change in C.M. exceeds a specified threshold) {
      Send flooding message
    }
  }
}

```

**FIGURE 6.3** Functionality of a DTMC in DICAf.

```

While (destination not reached yet) {
  If (entered the jurisdiction of a DTMC) {
    Communicate with the DTMC, download appropriate C.M. values
    Compute maximum permitted speed
    Set actual speed equal to maximum permitted speed
    Upload actual speed to DTMC
    Compute best route and execute travel
  }
  else {
    continue travel
  }
}

```

**FIGURE 6.4** Functionality of a vehicle.

fully automated, unmanned, and stand alone, much like the traffic signal boxes at street corners. The computing elements in the vehicles may consist of very low-cost microprocessors such as the Intel 8086, Motorola 6809, and so on. Similar to the traditional TMC, the DTMC computers of DICAf may be easily upgraded with more

powerful processors as they become available. DICAf is robust and reliable. Where one or more DTMCs fail, the remainder of the DICAf system will continue to function unlike a complete breakdown of the traditional system following the failure of the TMC. The vehicles within the jurisdiction of the failed DTMCs will be temporarily locked out of the C.M. values but will resume their normal activities as soon as they enter into the jurisdictions of other functional DTMCs.

### 6.3 MODELING DICAf ON AN ACCURATE, REALISTIC, AND PARALLEL PROCESSING TEST BED

The key contribution of DICAf consists in distributing the overall task of routing all vehicles through the network among all the entities in DICAf—vehicles and DTMCs. Thus, in reality, for a given DICAf system with  $V$  vehicles and  $D$  DTMCs, the total number of coordinating computing engines is  $(V + D)$ . To understand its performance and its dependence on different factors, DICAf is first modeled and then simulated on a parallel processing test bed that is constituted by a network of workstations configured as a loosely coupled parallel processor. The simulation coupled with the test bed virtually resembles a real implementation with one exception. To facilitate the simulation of a realistic system, that is, with a reasonable number of vehicles, while every DTMC is represented by a workstation, the vehicles are modeled as tasks and executed by the workstations underlying the DTMCs. When a vehicle is traveling along a highway segment, its computations are performed by the workstation underlying the DTMC that controls the segment. Every vehicle is represented through a process that migrates from one DTMC to the subsequent DTMC in the form of a message; the vehicle must be represented through a data structure where the fields represent the vehicle's key parameters. The fields include preferred speed, actual speed, origin, and destination. When a vehicle travels from the current DTMC (say A) to another DTMC (say B), the key parameters of the corresponding process in the underlying workstation for A is encapsulated through a message, propagated to B, and remanifested as a process in the underlying workstation at B. Thus, vehicles move in the simulation at electronic speeds instead of their physical speeds and, during its travel along a segment, a vehicle's computation subtask is executed on the workstation underlying the DTMC that controls the highway segment. Thus, DICAf is capable of simulating a highway system many many times faster than the actual operation.

This chapter assumes that a representative highway system is organized in the form of a rectangular grid, with any loss in generality. A DTMC is located at every intersection, and while it may be connected to a maximum of four segments in four directions it may only control a maximum of two segments—east and south. For DTMCs located at the bottom of the highway system, there are no south segments. Also, for DTMCs located to the extreme right of the highway system, there are no east segments. Thus, the maximum number of segments they control is unity.

Every DTMC maintains key information on the highway segments that it controls utilizing the following data structure: the first field, "index," stores the DTMC's

unique identifier, based on its location on the rectangular grid. While the second through fifth fields relate to the highway segment to its east, the sixth through ninth fields correspond to the highway segment to the south of the DTMC. For the east segment, the C.M. is stored in the field “eastCM” in mph units. Where the segment is nonexistent, this field is assigned a value of  $-1$ . The “eLastSent” stores the most recent value of C.M. for the segment that had been propagated to other DTMCs and “eSeq” refers to the sequence number of this value. The field “eCar” stores the number of vehicles currently traveling along the segment. The last four fields, southCM, sLastSent, sSeq, and sCar, correspond to eastCM, eLastSent, eSeq, and eCar respectively.

```
typedef struct DTMC {
    int     index;
    float   eastCM;
    float   eLastSent;
    int     eSeq;
    int     eCar;
    float   southCM;
    float   sLastSent;
    int     sSeq;
    int     sCar;
}
```

In addition, every DTMC maintains the C.M. values for all other highway segments in DICAF. These values are updated with new information received periodically from other DTMCs in the form of flooding messages. When a vehicle connects with a DTMC, it downloads a part of this database for determining its routing. For efficiency of memory usage, each DTMC declares an array of pointers of size MAX\_DTMC, as shown below. Additional memory is dynamically allocated during initialization to store the C.M. values of all highway segments corresponding to the specific highway system modeled in DICAF.

```
typedef struct DTMC *DTMC_PTR;
DTMC_PTR DTMC[MAX_DTMC];
```

DICAF utilizes two types of messages—one to encapsulate a vehicles and emulate the migration of the vehicles from one DTMC to the subsequent DTMC and another to propagate updated C.M. values for the highway segments. The data structure, shown next, is used to encapsulate vehicles. The “origin” field stores the DTMC identifiers where the vehicle originates while “dest” stores the final destination. The vehicle’s unique identifier is stored in “id.” The identifier is assigned at the DTMC where the vehicle is generated, and it is synthesized utilizing the DTMC identifier and the sequence number of the vehicle generated at that DTMC. The “update” and “from” fields store information to instruct the receiving DTMC to update local information

on the C.M. and other measures relative to the highway segment. The “update” flag is set if the current highway segment on which the vehicle is traveling is owned by the receiving DTMC. The “from” flag distinguishes whether the east or south segment of the DTMC requires updating. While the “inserted” field contains the time of assertion of the vehicle into a segment, the fields “appSec (in seconds)” and “appMSec (in milliseconds)” collectively store the time at which the vehicle is scheduled to appear, that is, arrive, at the destination DTMC. Evidently, this is computed from the length of the segment and the speed of the vehicle along it. The fields “trvSec (in seconds)” and “trvMsec (in milliseconds)” contain collectively the cumulative actual travel time of the vehicle from the origin to its current position while “idealTrvSec (in seconds)” and “idealTrvMsec (in milliseconds)” store the ideal time that the vehicle requires from origin to final destination, that is, the theoretical minimum travel time. While “speed” stores the speed of the vehicle along the current highway segment, the “prefSpeed” refers to its desired speed.

```
typedef struct Vehicle {
    int    origin;
    int    dest;
    int    id;
    int    update;
    int    from;
    long   inserted;
    long   appSec;
    long   appMSec;
    long   trvSec;
    long   trvMSec;
    long   idealTrvSec;
    long   idealTrvMSec;
    float  speed;
    float  prefSpeed;
}
```

The behavior of the workstation underlying every DTMC is presented in [Figure 6.5](#), in pseudocode.

In [Figure 6.5](#), the first two statements correspond to initialization. During initialization, the workstations underlying the DTMCs are interconnected through software links such that the resulting network corresponds to the highway system. In addition, vehicles are generated at the DTMC with unique identifiers and stochastic destinations. While their desired speeds are also stochastic, they conform to a normal distribution. The generated vehicles are included in the wait-to-leave list for the DTMC and are organized according to the times when they either enter the segment or depart the DTMC along an appropriate highway segment. The DICA simulation continues until the end of simulation time is reached. The simulation steps through the following phases. First, the incoming messages are read and processed by the function `Process_Message`, shown in [Figure 6.6](#). Where the message corresponds to

```

main {
  Synthesize the DICAF network;
  Generate vehicles and include in wait-to-leave list;
  While (simulation time is not expired) {
    Check incoming message(s) -- may be encapsulated vehicle or C.M. update;
    Process incoming message -- process vehicle or forward C.M. updates;
    Check wait-to-leave list;
    Process vehicle(s) off from the list;
    Propagate updated C.M. values;
  }
}

```

**FIGURE 6.5** The “main” routine in DICAF.

```

Process_Message {
  if (Message relates to flooding)
    Receive_Flooding;
  else {
    if (Vehicle commences travel of segment owned by current DTMC) {
      Update segment;
    }
    Insert vehicle into wait-to-leave list;
  }
}

Receive_Flooding {
  if (Message sequence number > previous sequence number of the segment) {
    Update segment -- update C.M. & advance sequence number;
    Forward flooding messages to all its neighbors;
  }
}

```

**FIGURE 6.6** The “process message” routine in DICAF.

a flooding message, that is, a C.M. update, it is used to update the local database and is then forwarded to other DTMCs. If the incoming message is an encapsulated vehicle, it is remanifested as a vehicle-process within the current DTMC and is inserted into the wait-to-leave list. Second, the wait-to-leave list is examined to check whether any vehicle must depart the DTMC at the current time. The vehicle’s routing is computed, encapsulated in the form of a message, and propagated along to the subsequent DTMC. Simultaneously, the C.M. values of the affected highway segments are recomputed and propagated to other DTMCs. The process continues until the simulation terminates.

Figure 6.7 details the organization of the function “check wait-to-leave list,” in pseudocode. During simulation, when the current time equals the time of departure from the DTMC of a vehicle, contained in the wait-to-leave list, the vehicle is first extracted from the list. Then, utilizing the most recent C.M. values of the highway

```

Process_vehicle {
  if (entry refers to a "shadow" vehicle) {
    Update segment since vehicle reaches the end;
  }
  else {
    if (current segment used by vehicle is owned by DTMC) {
      Update segment since vehicle exits the segment;
    }
    Determine best route and the subsequent DTMC;
    if (this DTMC is the final destination of vehicle) {
      Record relevant statistics in vehicle's data structure;
    }
    else if (subsequent segment is owned by current DTMC) {
      Determine the vehicle's speed on the segment;
      Update segment since vehicle enters the segment;
      Create a "shadow" entry for time when vehicle will complete travel
      Include entry in the list with shadow flag set
    }
    else if (current DTMC does not own highway segment) {
      set the "update" field in vehicle's structure
    }
    Propagate the data structure to the next DTMC;
  }
}

```

**FIGURE 6.7** "Check wait-to-leave list" function.

```

Update_segment {
  if (vehicle exits highway segment)
    Decrement number of vehicles on highway segment;
  else
    Increment the number of vehicles on highway segment;
  Compute new C.M. value;
  Update local copy of the C.M. value;
  if (C.M. differs from previous value by a margin exceeding threshold) {
    Propagate C.M. value through Send_flooding;
  }
}

Send_flooding {
  Increment sequence number associated with the segment;
  Initialize message with sequence number, segment ID, and C.M. value;
  Propagate message to all the neighbors;
}

```

**FIGURE 6.8** "Update\_segment" function in DICAf.

segments, the vehicle's routing is determined, and it is propagated to the subsequent DTMC along the appropriate segment. For a given segment, the value of C.M. may change when either a vehicle enters or exits. This causes the list to contain two kind of entries—(1) vehicles entering the segment that are labeled normal and (2) "shadow"

vehicles exiting the segment. Since only the owner DTMC maintains the C.M. for a segment, the “shadow” entry helps to ensure that the C.M. is updated when the vehicle completes its travel along the segment.

When the updated C.M. value for a highway segment differs substantially from its previous value, the owner DTMC must propagate the information to other DTMCs through the flooding scheme. The responsible function is “Update\_segment.” In this scheme, a sequence number is associated with every message to cut down on duplicate transmissions. When a message is broadcast from an originating DTMC to its neighbors, the DTMC increments the sequence number. When this message is received by a different DTMC, it will forward the message to its neighbors only if the sequence number is more recent than that associated with the previous copy of the C.M. The algorithm is presented in pseudocode in [Figure 6.8](#).

## 6.4 IMPLEMENTATION AND DEBUGGING ISSUES

The DICAF model and simulator is written in C and designed to execute on a network of Unix-based workstations, connected through a 10 Mbit/s Ethernet and configured as a loosely coupled parallel processor. The network includes over 65 SUN Sparc 10 workstations under SUN Solaris 2.4 operating system. The code segment for every DTMC including the vehicles introduced into the system at that location is approximately 4432 lines of C code while the networking code is approximately 1500 lines of C++ code. The simulator is compiled by the public domain GNU C compiler, gcc and is executed at a low priority in the background utilizing the “nice” utility. The execution of DICAF generates approximately 4 to 9.5 MB of data, which is then parsed by a parser to yield the performance graphs. The parser consists of 910 lines of C. Each simulation run corresponds to an experiment, which, in turn, represents 24 hours of actual highway traffic and requires approximately 16 minutes of wall clock time for both a 10 processor and 51 processor DICAF system. The maximum number of vehicles introduced into DICAF is 45,000. The underlying test bed consists of 65 concurrently executing SUN Sparc 10 workstations, each with 32 MB of RAM and a 424-MB hard drive. A total of 400 simulation runs are executed throughout this research. It may be noted that the workstations may be executing primary jobs for the users at the respective consoles.

A unique characteristic of asynchronous, distributed algorithms, DICAF being a specific instance, is that they are extremely complex to debug. It is estimated that debugging a 50-node DICAF simulation is equivalent to debugging approximately  $4,000 \times 50 = 200,000$  lines of C code, given that DICAF is approximately 4,000 lines long. In truth, however, the problem is more severe. At any time instant during execution, every processor may be executing a unique line of the DICAF program and there may be little to no correlation between them. Since we human beings think sequentially, in general, comprehending the simultaneous execution of multiple autonomous entities is very difficult, at the least. Conceivably, one can add printf statements in the code to display the values of key variables and structures as the execution progresses. However, the numbers of times that the values of the variables



are printed out quickly runs into the thousands and becomes overwhelming. Timing-related problems are more acute and even harder to debug since they can become intermittent if `printf` statements are added to the code. In the experience of the authors, efficient debugging requires an extremely thorough knowledge of the program, very high degrees of concentration, commitment, and patience.

Asynchronous, distributed algorithms also impose great demand on the accuracy and correctness of the underlying operating system. First, such algorithms require fast propagation of small to modest size messages. Second, the distribution of messages is highly bursty. Third, given that each processor must execute the route guidance computations for 45,000 vehicles in addition to computing and disseminating the congestion measures, the burstiness of the message distribution reaches an extreme. This, in turn, imposes frequent writes into the buffers that are constantly full. It was observed that at a few times the simulation fails unexpectedly, accompanied by the errors: can't write, broken pipe, and sigpipe error. Other, hard-to-explain errors include a scenario where one of the processors fails to write data into a file, despite successfully opening it.

During simulation, a vehicle collects and stores its travel-related information including routing, C.M. values, and time of travel. The data is written out into a file when the vehicle reaches the destination DTMC. The parser utilizes it to generate performance statistics including the average speeds of vehicles throughout their travel, the actual elapse times, the theoretical minimum travel times, the actual speed distributions of the vehicles, the percentage differential travel times as a function of the vehicle identifiers and as a function of the time of assertion of the vehicle into DICAF. The percentage differential travel time is defined later in this chapter. In addition, the parser assists in plotting the C.M. on the highway segments, both as functions of the highway segments and the simulation time. The distribution of the number of vehicles on the highway segments and the number of vehicle-DTMC and inter-DTMC messages are also obtained by the parser. Corresponding to a number of measures, the parser also computes the average and standard deviation values.

## 6.5 SIMULATION RESULTS AND PERFORMANCE ANALYSIS OF DICAF

The ultimate objectives of DICAF include (1) fast arrival of all vehicles at their destinations and (2) superior utilization of all resources, that is, highway segments. In addition, DICAF aims to deliver a scalable, pragmatic, yet economical system wherein the communication and computational demands on the individual DTMCs are minimal. Furthermore, for a given highway system and a given realistic, input traffic distribution over a time interval of interest, the travel time required for any vehicle between any given origin and destination, in DICAF, must be uniform and independent of the time at which the vehicle commences its travel during the interval.

For a systematic presentation of the performance of DICAF, in this chapter, first the independent parameters and key performance measures are identified and a number of simulation experiments are executed, corresponding to representative traffic

conditions, on two highway systems, shown in Figures 6.1 and 6.9. While the highway system in Figure 6.1 consists of 9 DTMCs and 12 highway segments, Figure 6.9 consists of 50 DTMCs and 85 highway segments. Every DTMC controls the segments to its east and south. All segments are assumed to be 10 miles in length, without any loss




FIGURE 6.9 A second representative highway system modeled in DICAF.

in generality, since the speeds of the vehicles, as described later, are stochastically determined.

The independent parameters include (1) the number of vehicles asserted into the system, (2) the density of vehicles, that is, the frequency with which the vehicles are input into DICAf, and (3) the distribution of vehicle speeds.

The key performance measures include (1) the distribution of the travel times of vehicles arriving at their destinations, (2) the fraction of vehicles reaching their destinations, (3) the average and standard deviation of the travel times as a function of traffic density, (4) the distribution of the C.M. for selected highway segments as a function of the simulation time, (5) the average C.M. values for the highway segments for different traffic densities, (6) the travel times of the individual vehicles as a function of the times of their assertion into DICAf, and (7) the average number of vehicles on each of the highway segments, under different traffic densities. In addition, two measures are presented to permit closer examination of DICAf's performance relative to shorter travel times of vehicles. The first (8) is the distribution of the actual average speeds of every vehicle during its travel, relative to its desired speed. The second measure aims to provide an insight into the impact of the distribution of desired speeds of vehicles, on DICAf. This is achieved by measuring the average of the travel times of vehicles for higher  $\mu$  value and identical  $\sigma$  value and for higher  $\mu$  value coupled with smaller  $\sigma$  value. These are labeled measures (9) and (10) respectively.

To estimate DICAf's achievement in distributing the overall communication task, two additional performance measures are defined. They include (11) the maximum data transfer rate from each of the DTMCs to the vehicles, measured every second of the simulation execution and (12) the distribution of the total number of flooding messages, both received and sent by the DTMCs, as a function of simulation time. The flooding messages carry the C.M. values for other highway segments from other DTMCs.

In DICAf, a vehicle is represented through a data structure and is propagated from one DTMC to another at electronic speeds, unlike the slow actual highway speeds. Thus, the DICAf simulation can execute many times faster than an actual highway operation, and this permits a detailed study of DICAf under many representative scenarios. The DICAf design is organized to execute 90 times faster than an actual highway operation. That is, 1 second of the simulation corresponds to 90 seconds of actual highway operation and a 24-hour highway operation is simulated in  $(24 \times 60 \times 60 \div 90) = 960$  seconds or 16 minutes of wall clock time. This is achieved in DICAf as follows: in DICAf, a vehicle arrives at a DTMC from the preceding DTMC in only a few milliseconds. However, the arrival of the vehicle is deliberately delayed at the destination DTMC by a value that corresponds to the travel time in the actual highway divided by 90. This also guarantees the accuracy of modeling and representing the travel of every vehicle in DICAf relative to the resolution of the simulation.

The average vehicle length is assumed to be 0.0025 mile, which yields the maximum number of vehicles allowed on any segment at  $0.5 \times \frac{10}{0.0025} = 2000$ , as per Equation 6.1. The input traffic distribution uses a stochastic function to generate vehicles that are uniformly asserted at the DTMCs of DICAf. Every simulation

run executes corresponding to 24 hours of actual highway operation. Accordingly, the traffic generator asserts vehicles during the simulation corresponding to a 24-hour day. At every DTMC, the vehicles generated are assigned destination DTMCs stochastically. The number of vehicles asserted at any DTMC is a uniform function of time throughout the simulation. However, the vehicles' speeds are stochastically generated, utilizing the "drand48" pseudorandom generator function, and they follow a normal distribution with specified  $\sigma$  and  $\mu$  values. Every vehicle is assigned a unique identifier. For the 9-DTMC highway system, the identifier is computed as:  $([\text{DTMC identifier} \times 10,000] + [1 \dots \text{maximum number of vehicles generated at that DTMC}])$ . For the 50-DTMC highway system, the identifier is computed as:  $([\text{DTMC identifier} \times 1,000,000] + [1 \dots \text{maximum number of vehicles generated at that DTMC}])$ . The maximum C.M. value allowed for the highway segments is 100 mph, and the highest and lowest permitted vehicle speeds are 90 mph and 30 mph respectively. Thus, in this chapter, most of the performance results for both highways systems in Figures 6.1 and 6.9 are obtained for  $\mu = 50$  mph and  $\sigma = 10$  mph. In addition, the highway system of Figure 6.1 is simulated in DICAF for the scenarios—(1)  $\mu = 65$  mph and  $\sigma = 10$  mph and (2)  $\mu = 65$  mph and  $\sigma = 5$  mph. While the highway system in Figure 6.1 requires 10 processors that execute concurrently and asynchronously, Figure 6.9 requires 51 simultaneously executing SUN Sparc 10s. The number of vehicles asserted into DICAF is controlled by the traffic density function that assumes the values 3, 6, and 8, for the highway system in Figure 6.1. A traffic density value,  $D$ , implies that, at every DTMC, a stochastically generated number of vehicles, ranging from 0 to  $D$ , is asserted into DICAF at every second of the simulation. The corresponding numbers of vehicles are 8,528, 21,572, and 30,284, respectively.

The RIDOT maintains traffic density for the key highways including I-95, I-195, and I-295, in terms of the maximum number of vehicles carried by each of the segments over a 24-hour period. Since the same vehicle may travel on different highway segments, it is nontrivial to estimate the total number of vehicles asserted into the entire Rhode Island highway system from the traffic density map [89]. It is noted that the 3-lane, I-95 highway carries the most traffic, namely a maximum of 89,302 vehicles throughout a 24-hour period. Therefore, the maximum, normalized traffic density, relative to a single-lane highway, will correspond to 30,000 vehicles on a 24-hour period. Given that Rhode Island is a tiny state, approximately 40 miles by 16 miles in size, with the I-95 assuming the role of the most important roadway, this chapter assumes that the I-95 constitutes a part of the itinerary of nearly every vehicle that travels on any highway segment of Rhode Island. Therefore, the maximum number of vehicles asserted into the Rhode Island highway system is approximately 30,000 over a 24-hour period. Thus, the highest traffic density utilized in DICAF, namely eight, is representative of actual traffic conditions.


For the 50-DTMC highway system shown in Figure 6.9, the total number of vehicles asserted into DICAF is 45,296, corresponding to a density of 3. The simulation is also designed to execute 90 times faster than the actual highway operation.

It is noted that each vehicle is generated stochastically, that is, its source, destination, and desired speed of travel are all stochastic quantities. Given that more

than 45,000 vehicles are generated and simulated, the results obtained are assumed, reasonably and justifiably, to yield a general insight into the performance of DICAF. Although more than 400 simulation runs were executed, each time with different seed values, the general behavior of the data was observed to be similar, and the results reported here correspond to a specific run. It may be further noted that neither  $\mu$  nor  $\sigma$  values were measured. They were assumed, for the sake of generality. The confidence interval for a  $\mu$  is given by  $\{\hat{x} - k \leq \mu \leq \hat{x} + k\}$ , where  $\hat{x}$  is the computed mean and  $k = \frac{c\mu}{\sqrt{n}}$ , where  $c$  relates to the confidence level  $\gamma$ ,  $\sigma$  is the measured variance, and  $n$  is the number of samples. In this chapter,  $n$  is over 45,000 which implies a very small value for  $k$ , that is, a very narrow confidence interval, even for a large value of  $c$ , that is, very high confidence level.

The performance data, presented in this section, is obtained from sampling key measures every simulation second.

Figure 6.10 presents the normal distribution of the desired speeds of the vehicles for traffic density values of 3, 6, and 8, for the 9-DTMC highway system. The x-axis represents the speeds of the vehicles, ranging from 30 mph to 90 mph while

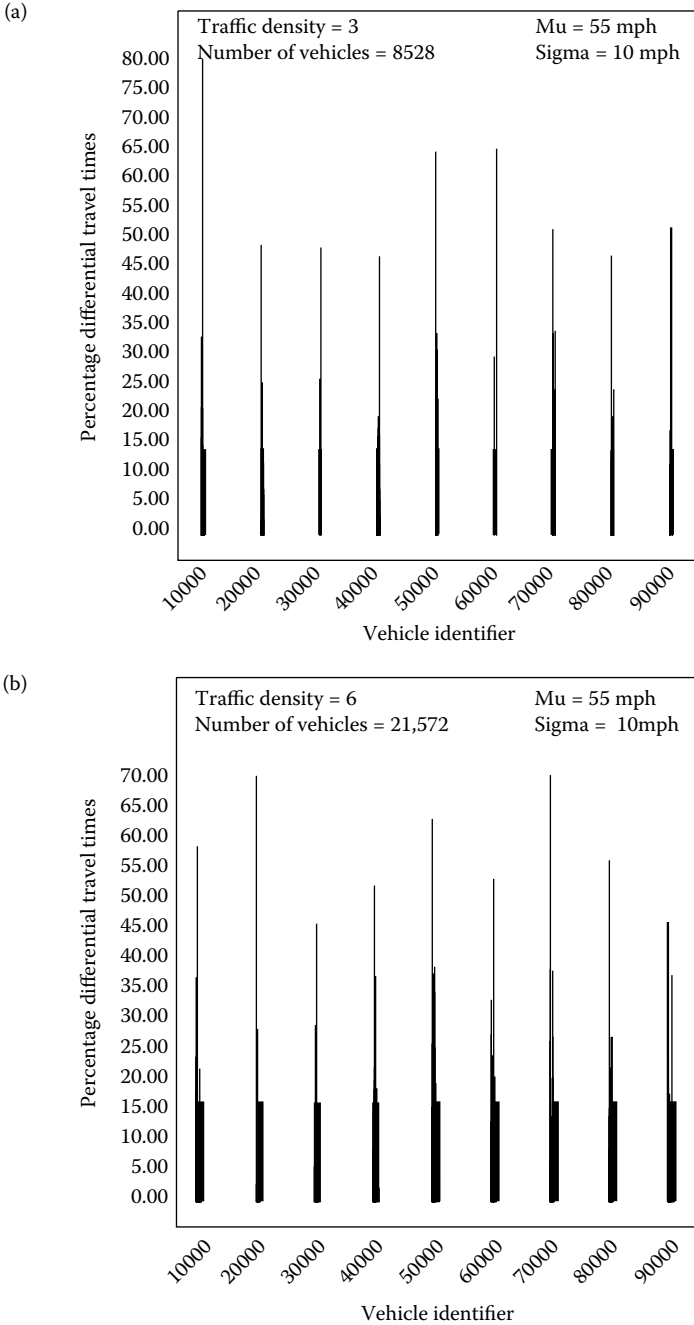


**FIGURE 6.10** Normal distribution of the desired speeds of vehicles for traffic density values of 3, 6, and 8.

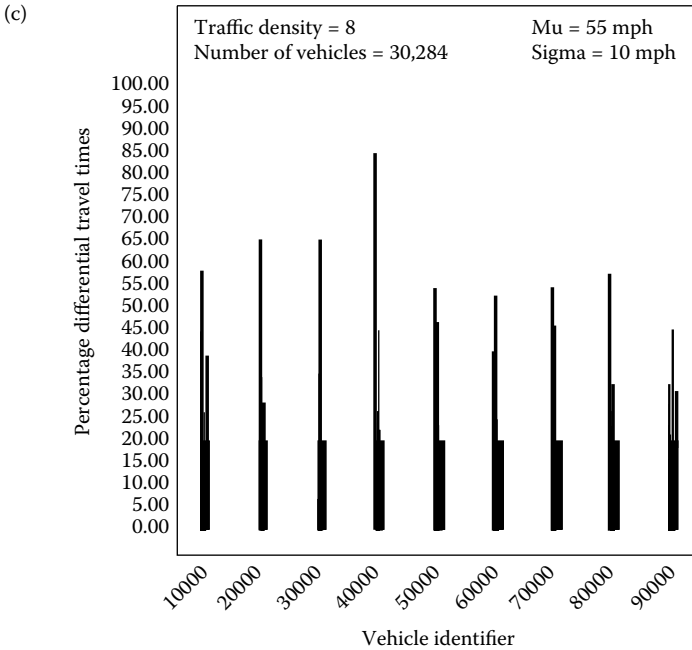
the y-axis presents the number of vehicles corresponding to the respective desired speeds. Clearly, the bulk of the vehicles desire travel speeds around 55 mph while fewer vehicles desire to travel at the lower and higher speeds.

To contrast the performance of DICAF against the traditional TMC scheme requires the development of a simulator that utilizes centralized scheduling and executes sequentially on a single processor. In this research, a uniprocessor simulator is not developed for two reasons. First, the memory requirement to represent the data structures corresponding to an excess of 30,000 vehicles is prohibitively large. Second, the speed of execution of the simulator is likely to be excruciatingly slow. This chapter proposes to evaluate the absolute performance of DICAF. In this chapter, corresponding to every vehicle, the time to complete the travel is computed as the length of the shortest path from the origin DTMC to the destination DTMC, in miles, divided by the desired speed of the vehicle. This corresponds to the ideal travel time since it may be achieved by the vehicle only in the complete absence of any other vehicle in DICAF competing with it for resources. The ideal travel time for a vehicle is the best that it can achieve, in the absolute sense. When a simulation run is complete, the actual travel time required by each of the vehicles is obtained from the information that the vehicles collect as they progress from their origins to the respective destinations. For a vehicle, the actual travel time is the cumulative sum of the travel times between every DTMC pair in its route during the simulation. The percentage differential travel time for a vehicle is then computed as:  $([\text{actual travel time for a vehicle} - \text{ideal travel time}] / [\text{ideal travel time}] \times 100)$ . By definition, the percentage differential travel time for every vehicle must be a positive percentage, that is, greater than or equal to 0%, and it reflects the travel time that the corresponding vehicle requires in excess of the absolute minimum. Figures 6.11a through 6.11c plot the percentage differential travel time along the y-axis for every vehicle in DICAF, for density values 3, 6, and 8, respectively. The x-axis represents the unique vehicle identifiers that range from 10,000 to 90,000. The average and standard deviation value pairs for density values 3, 6, and 8 are {1.89%, 6.09%}, {1.85%, 5.71%}, and {2.12%, 6.14%} respectively. The increase in the average of the percentage differential travel times overall vehicles, from 1.89% to 2.12%, corresponding to traffic densities three and eight respectively, reflects the increased competition for highway segments by more vehicles, which, as expected, results in greater travel times for all vehicles. Furthermore, with more than 30,000 vehicles asserted into DICAF, the average vehicle's travel time is only 2.12% higher than the absolute best. In addition, the worst-case travel time, shown in Figure 6.11c, is only 85% more than the absolute best or less than twice the ideal travel time.

Clearly, these findings attest to DICAF's strong performance. Thus, despite 30,284 vehicles competing for highway segments, the autonomous, distributed, and dynamic routing of each vehicle in DICAF yields results that are very close to ideal. The graphs are especially revealing for the following reason: in the parallel processing community, while it is generally accepted that the use of distributed algorithms may yield faster results, the issue of the quality of the solution relative to that from the centralized algorithm is an open one. In distributed algorithms, local entities execute



**FIGURE 6.11** Percentage differential travel times for all vehicles [calculated as (actual travel time for each vehicle – ideal travel time) / (ideal travel time) × 100%] for (a) density 3, (b) density 6, and (c) density 8.



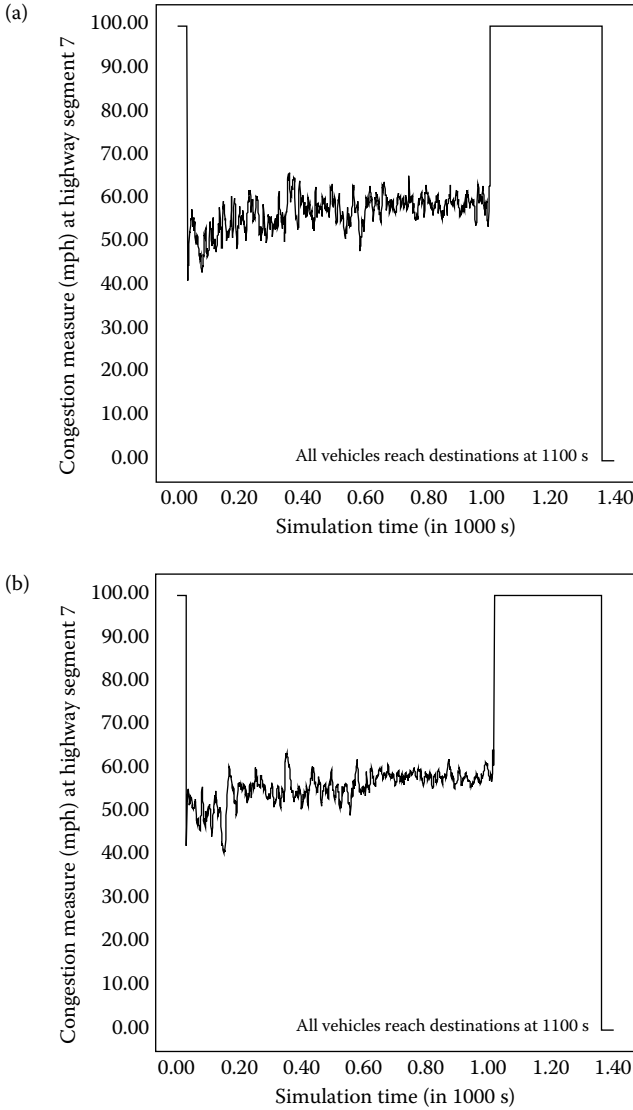
**FIGURE 6.11** (continued)

the decision making, while allowed access to only a fraction of the system-wide state. While the lack of access to the system-wide state may raise concern, the results of DICAF reveal that the quality of the results is very high. That is, as evident from DICAF, under certain circumstances, distributed algorithms may yield very high-quality solutions while also generating them fast. The authors have developed [90] a new mathematical framework to extract distributed algorithms from centralized descriptions of problems.

In the DICAF simulation, vehicles are uniformly asserted at the DTMCs throughout the entire simulation run, that is, up to 960 simulation seconds. The simulation is permitted to continue execution up to 1100 seconds when it is observed that all vehicles arrive at their respective destinations.

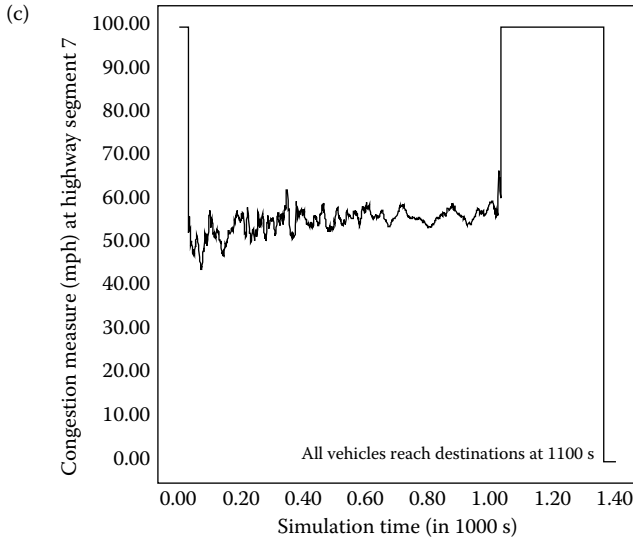
Figures 6.12a through 6.12c present the variation of the C.M. value for a selected highway segment, 7, under the traffic density values 3, 6, and 8, respectively. The variation of the C.M. value is continuous and consistent with the definition in Equation 6.1. The reason for the choice of segment 7 is that it is located at the center of the highway system and many vehicles traveling diagonally are likely to include it in their itinerary. The initial, default C.M. value is 100 mph while, following the completion of travel of all vehicles, that is, at 1100 simulation seconds, the C.M. value increases to 100 mph. As expected, the average C.M. value is higher for low traffic density and lower when the total number of vehicles is high.





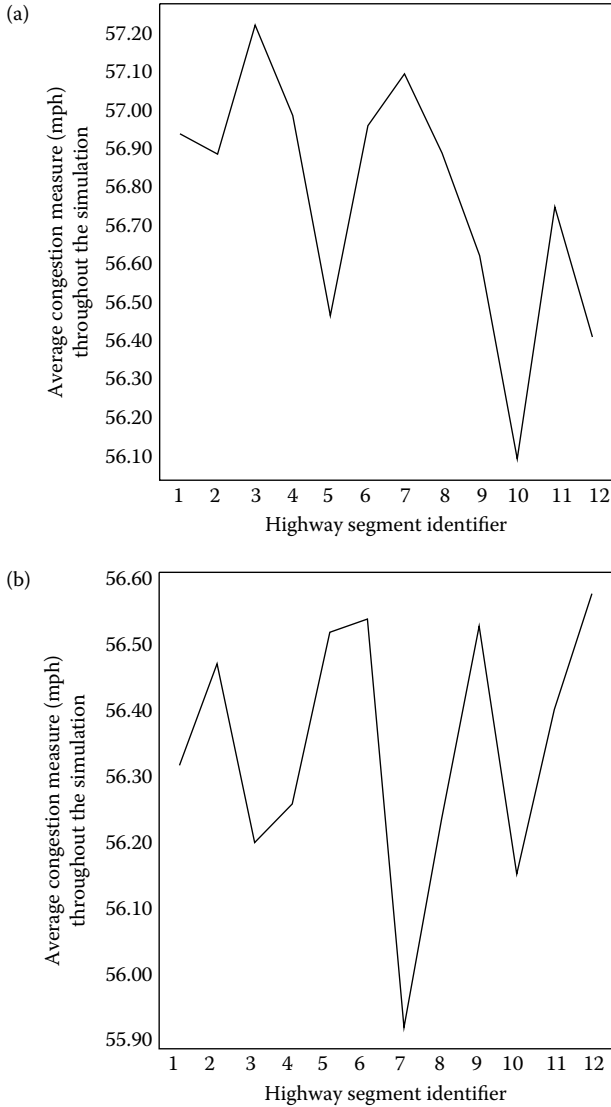
**FIGURE 6.12** Distribution of C.M. for highway segment 7 as a function of the simulation time for (a) density 3, (b) density 6, (c) density 8.

Figures 6.13a through 6.13c present the average C.M. values for all highway segments, computed over the entire simulation run, for each of the three traffic density values. It is noted that the average C.M. values for the highway segments are close to one another, ranging from 54 mph to 57 mph, implying that all resources are equitably utilized in DICAf.



**FIGURE 6.12** (continued)

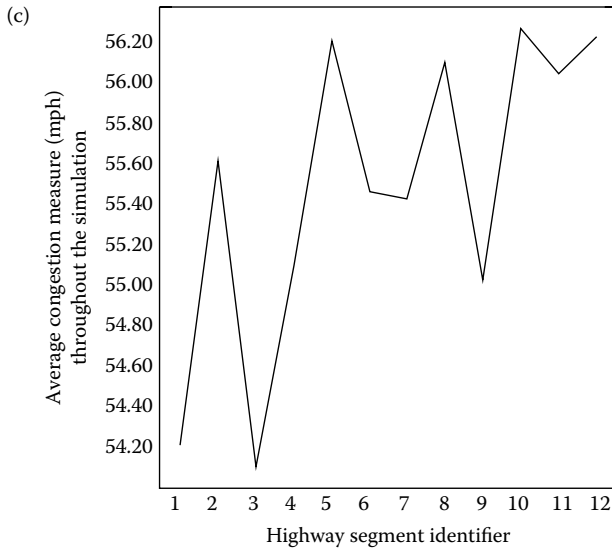
Figures 6.14a through 6.14c present the highway resource utilization through a plot of the average number of vehicles, over the entire simulation run, for each of the highway segments, for the three traffic density values. Although all highway segments are utilized reasonably, which reflects DICAF’s success in utilizing all highway resources, the average number of vehicles is higher on some segments more than others, principally due to the stochastic speeds and destinations of the vehicles. The average number of vehicles for every segment is higher in Figure 6.14c than in Figure 6.14b, which, in turn, is greater than in Figure 6.14a. This is expected since higher traffic density value implies a greater number of vehicles that must utilize the highway segments to achieve their travels. For density 3, the maximum of the average number of vehicles overall segments is 16, while for density 6 and 8 the corresponding figures are 83 and 130 respectively. The average number of vehicles for any segment is obtained by sampling the vehicle queue length, associated with the segment, every simulation second. To relate these figures to the maximum capacity of the highway segments, consider the following. A highway segment may support a maximum of 2000 vehicles on it, all of which must be traveling at the same speed anywhere from 0 mph to 90 mph. The shortest time for which this statement may be true, may be obtained as  $10 \text{ miles} \div 90 \text{ mph} = 399.9$  actual highway seconds, assuming an average speed of 90 mph. This, in turn, corresponds to 4.4 simulation seconds. Therefore, corresponding to 1 simulation second, the maximum number of vehicles on a highway segment is  $2000 \div 4.4 = 454$ . Thus, the maximum of the average number of vehicles on any segment under DICAF, namely 130, corresponds to approximately 28% of the absolute maximum. Thus, for a given number of vehicles, unlike the traditional approach where only a few segments are



**FIGURE 6.13** Average C.M. value for all highway segments of DICAf for (a) density 3, (b) density 6, (c) density 8.

extremely congested, traffic is equitably distributed among all highway segments in DICAf.

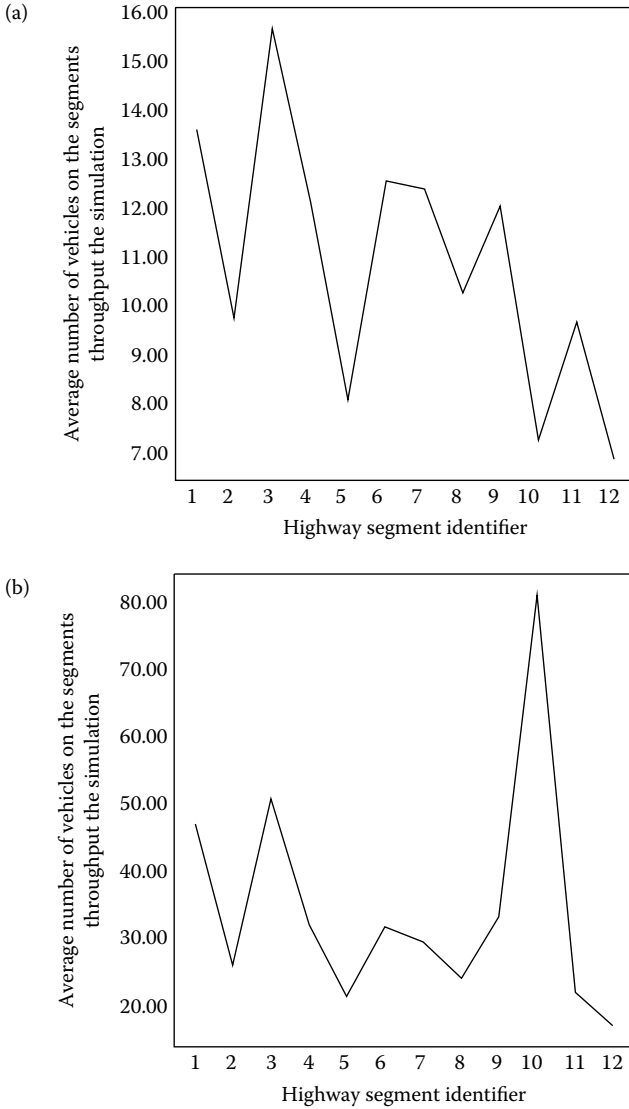
To understand the influence of the assertion time of the vehicles on their travel times, [Figures 6.15a](#) through [6.15c](#) present the tuples ([difference of actual travel time of a vehicle relative to its ideal travel time as a fraction of its ideal travel time], time



**FIGURE 6.13** (continued)

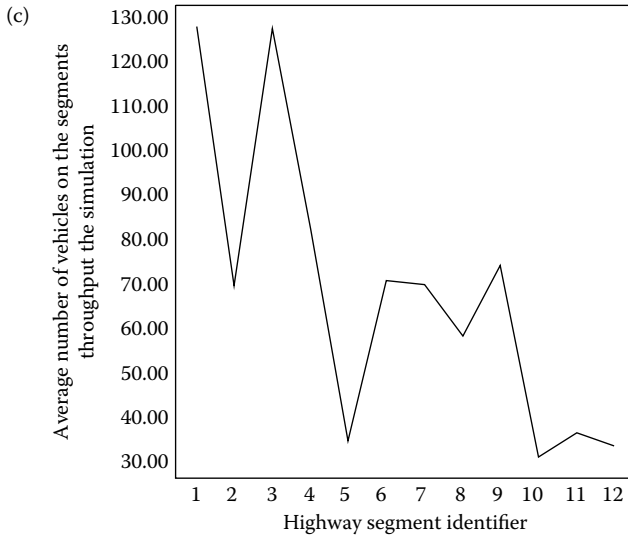
of assertion of the vehicle into DICAF) for all vehicles. Figures 6.15a through 6.15c correspond to traffic density values of 3, 6, and 8 respectively. The increased darkening of the graphs in Figures 6.15a through 6.15c reflects the increased number of vehicles. Except for a dark band between 600 seconds and 960 seconds in Figure 6.15c, the outlines of the three graphs are similar. This implies that, except at the beginning of simulation, vehicle travel times remain unchanged regardless of where, in the simulation time line, they are introduced into DICAF. The dark band between 600 seconds and 960 seconds in Figure 6.15c is not significant. It merely reflects the introduction of an appreciable number of vehicles into DICAF within the simulation time interval—{600 s, 960 s}—and that for these vehicles, the percentage differential travel times range from 0% to 10%. However, the negative slopes of the outlines of the three graphs, or stated differently, the absence of data points in the upper right hand corner of the graphs, is significant. It implies that, despite a realistic and high traffic density, DICAF successfully achieves efficient routing for all vehicles. Given the observation that in every simulation all vehicles reach their respective destinations correctly, the graphs in Figures 6.15a through 6.15c also imply that the given highway system is stable through the highest traffic density of 8, that is, it can accommodate at least 30,284 vehicles.

In the graphs, vehicles asserted into DICAF early in the day appear to take longer travel times relative to those that are asserted at all other times during the simulation. This is counterintuitive as one would normally expect very little to nil vehicle build up, due to congestion, early in the day. The reasoning is obtained from examining the speed distribution of the vehicles asserted into DICAF as a function of the simulation time, which is shown in [Figure 6.16](#).



**FIGURE 6.14** Average number of vehicles on the highway segments computed over the entire simulation run for (a) density 3, (b) density 6, (c) density 8.

Figure 6.16 reveals that the speeds of the vehicles, asserted into DICAf, range from 30 mph to 90 mph from simulation time 0 to 100 seconds. From 100 to 550 seconds, the speed range narrows to a band—{40 mph, 80 mph}—which is followed by an even narrow band—{50 mph, 70 mph}—up to 960 seconds. Given the goal that the vehicle speeds must follow a normal distribution, as shown in Figure 6.10, the



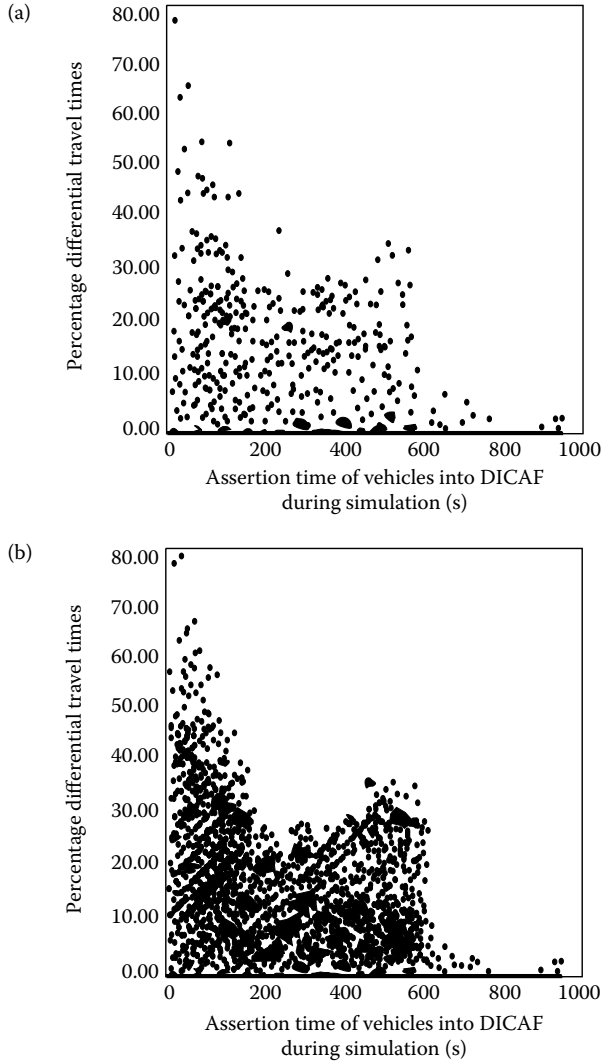
**FIGURE 6.14** (continued)

traffic generator program first creates “speed buckets,” with speed ranges extending from {30 mph, 35 mph} to {85 mph, 90 mph}, and then allocates the appropriate number of vehicles for each bucket and for each of the three traffic density values. Next, vehicles with stochastic destinations and speeds are generated, and they are assigned to the appropriate buckets to meet the appropriate allocation, starting at simulation time 0 seconds. As expected, the buckets with the extreme ranges—{30 mph, 35 mph} and {85 mph, 90 mph}—get filled first, that is, at lower values of simulation time. Thereafter, the buckets with the subsequent extreme speed bands—{40 mph, 45 mph}, {45 mph, 50 mph}, {70 mph, 75 mph}, and {75 mph, 80 mph}—are filled at simulation times ranging from 100 to 550 seconds. By this time, all buckets except those in the range {50 mph, 70 mph} have been filled. So, from 550 to 960 simulation seconds, the vehicles asserted into DICAF bear speeds within a narrow band.

Logically, the slower speed vehicles are very likely to lower the C.M. values of the segments. Therefore, in addition to the expected longer travel times of the slow vehicles, those with higher desired speeds will also require longer travel times. As simulation progresses, the average desired speeds of the vehicles is observed to increase, resulting in faster travel and shorter travel times.

In addition, a comparison of the graphs in [Figures 6.17](#) and [6.10](#) reveals that while the actual speed distribution of vehicles resemble the initial desired normal distribution, the distribution in [Figure 6.17](#) is somewhat flattened due to congestion.

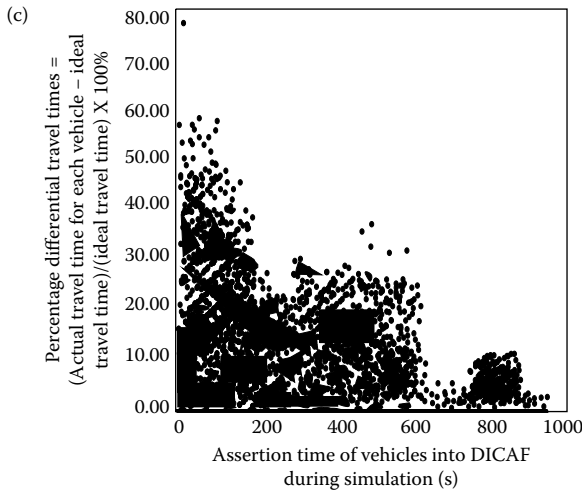
[Figure 6.17](#) plots the actual average speeds of every vehicle, obtained through dividing the total distance traveled by each vehicle by its respective travel time. It may be observed that, throughout the simulation, vehicles travel at average speeds



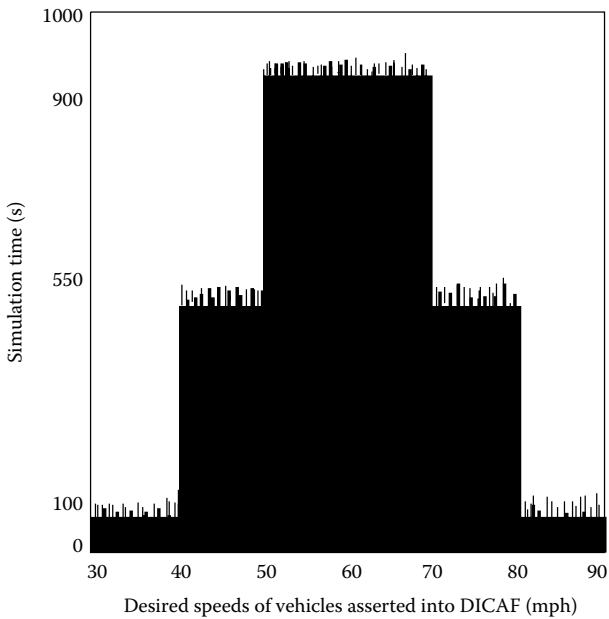
**FIGURE 6.15** Percentage differential travel times for all vehicles [calculated as (actual travel time for each vehicle – ideal travel time) / (ideal travel time) × 100%] as a function of the assertion time of the vehicle for (a) density 3, (b) density 6, (c) density 8.

ranging from 30 mph to 80 mph although the number of vehicles beyond 75 mph is minimal. The average speeds are crowded around 60 mph, implying that, despite competition for resources from over 30,000 vehicles, most of the vehicles are able to determine their routes efficiently in DICA F, thereby realizing shorter travel times.

The contrast between the percentage differential travel times for vehicles introduced into DICA F at the beginning of simulation versus those that are asserted later




**FIGURE 6.15** (continued)



**FIGURE 6.16** Desired speeds of vehicles asserted into DICAF for traffic density 8 as a function of the simulation time.

in the simulation, as observed in Figure 6.15, appears to indicate that narrow speed bands may bear significant impact on routing efficiency and faster travel times for all vehicles. To examine this hypothesis, simulations are designed to execute for two



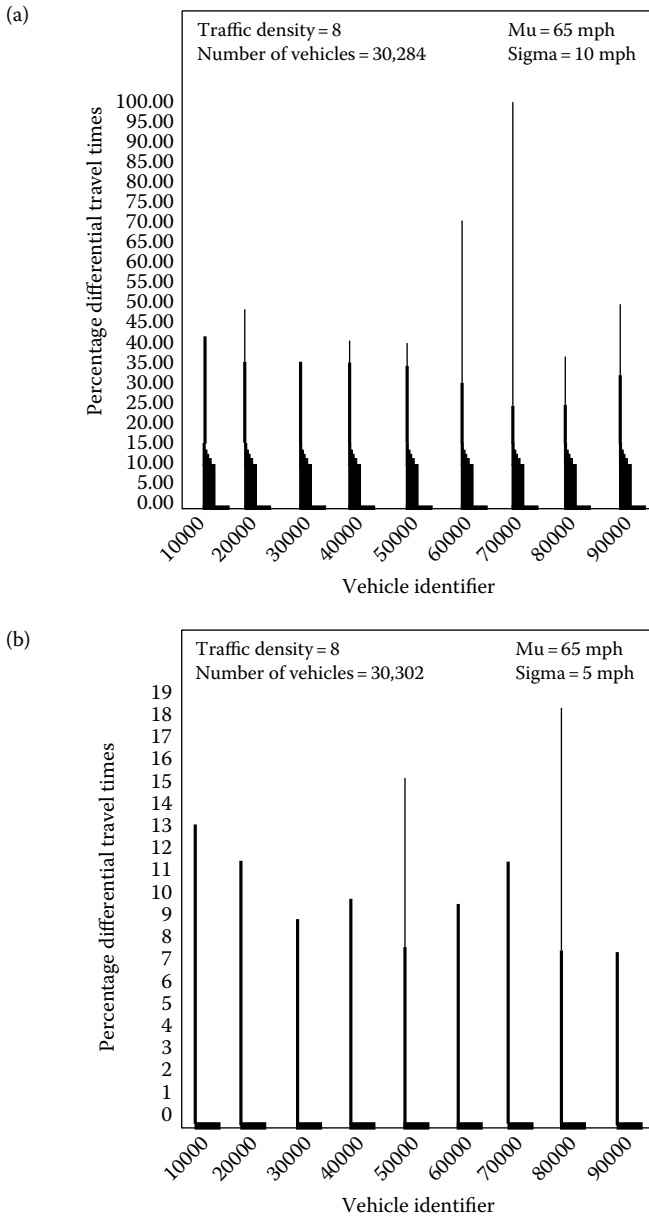


**FIGURE 6.17** Actual average speeds of vehicles in DICAF for traffic density 8.

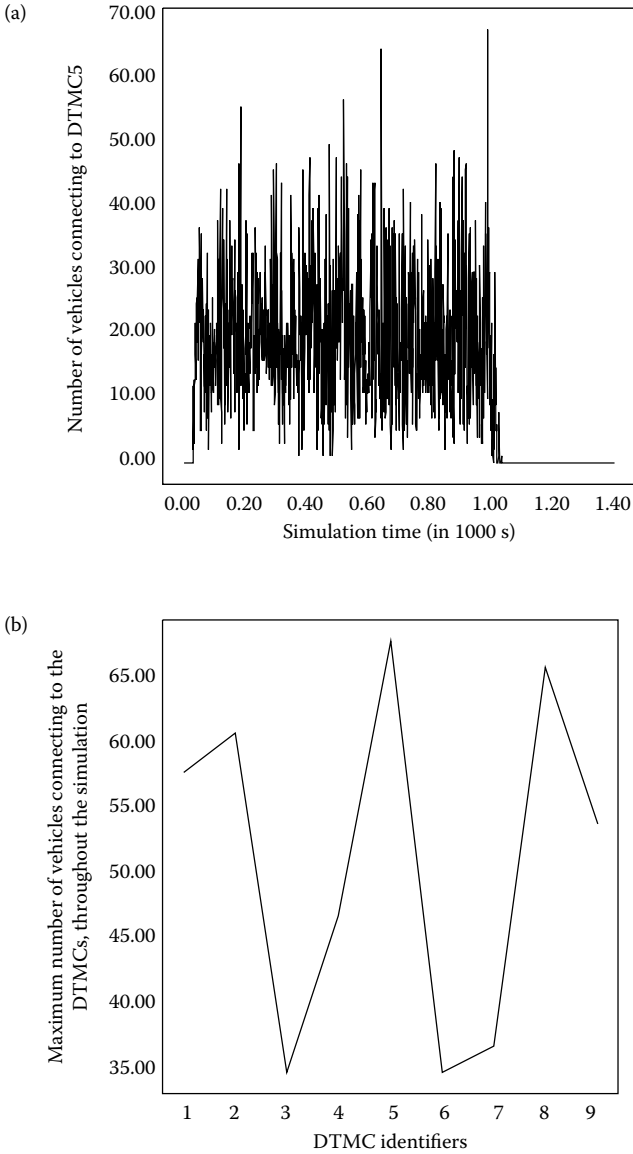
input traffic patterns— $\{\mu = 65 \text{ mph}, \sigma = 10 \text{ mph}\}$  and  $\{\mu = 65 \text{ mph}, \sigma = 5 \text{ mph}\}$ . The traffic density selected is eight. Figures 6.18a and 6.18b present the difference of actual travel time of each vehicle relative to its ideal travel time as a fraction of the ideal travel time for all vehicles. Figures 6.18a and 6.18b correspond to the traffic generator parameters— $\{\mu = 65 \text{ mph}, \sigma = 10 \text{ mph}\}$ . Corresponding to Figure 6.18a, the number of vehicles is 30,284 and the average and standard deviation values are 1.11% and 4.13%. For Figure 6.18b, the number of vehicles is 30,302 and the average and standard deviation values are 0.03% and 0.45%.

A comparison with the corresponding figures for Figure 6.11c reveals that travel efficiency, that is, shorter travel times for all vehicles, may be achieved with relative ease through mandating and enforcing a narrow speed band, defined by minimum and maximum permitted speed values, unlike the current policy of a single 55 mph maximum speed limit. Conceivably, in the current highway system, a key cause of speed fluctuation of vehicles arises from the lack of information on the highway conditions, such as exit location and speeds, all of which may be successfully resolved by DICAF.

A principal objective of DICAF is to equitably distribute the overall communications task to all entities. Figure 6.19a presents the distribution of the number of vehicles connecting to DTMC5 as a function of the simulation time. As explained



**FIGURE 6.18** Percentage differential travel times for all vehicles [calculated as (actual travel time for each vehicle – ideal travel time)/(ideal travel time) × 100%] for (a) { $\mu = 65$  mph,  $\sigma = 10$  mph}, (b) { $\mu = 65$  mph,  $\sigma = 5$  mph}.



**FIGURE 6.19** (a) Number of vehicles connecting to DTMC5 as a function of simulation time, (b) maximum number of vehicles connecting to the DTMCs.

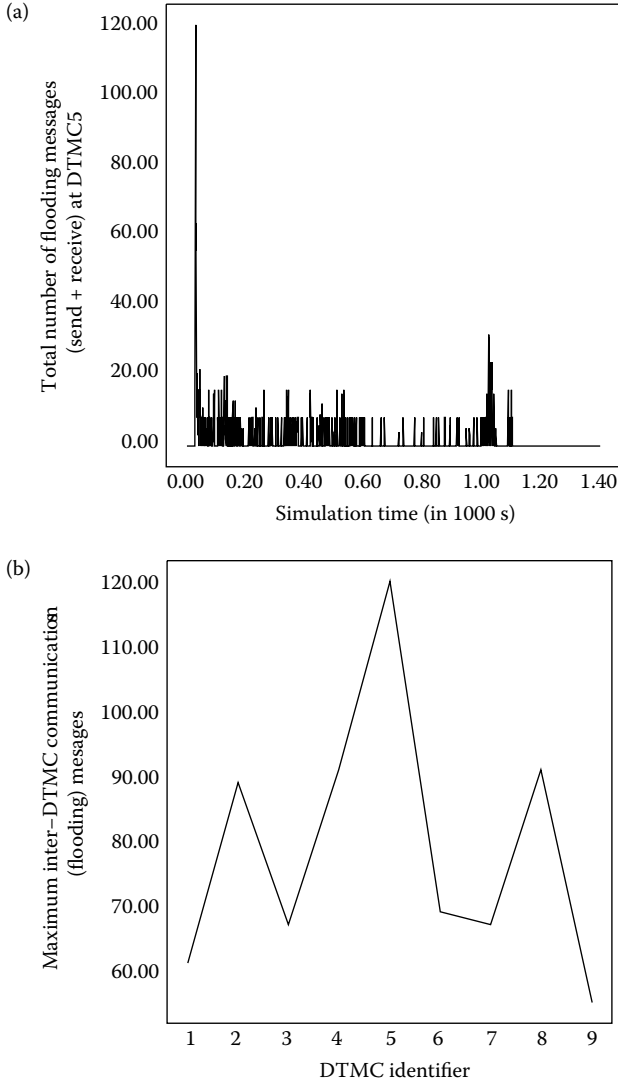
earlier, the reason for the choice of DTMC5 is that it is located at the center of the highway system, it is likely to be in the itinerary of many vehicles, and, as a result, the communications related data reflects the worst-case scenario. The graph in Figure 6.19a also reflects the fact that the number of vehicles on any highway

segment and the consequent C.M. of the segment are highly dynamic function of time. Figure 6.19b plots the maximum number of vehicles connecting to all of the DTMCs. The data presented in both Figures 6.19a and 6.19b are obtained from sampling the appropriate queues every simulation second. When a vehicle connects to a DTMC, it downloads the C.M. values of the appropriate segments and the average information is 96 bytes. Thus, corresponding to the maximum number of vehicles of 68 for DTMC5, the maximum DTMC-vehicle communications rate is given by  $68 \times 96$  bytes/simulation-second = 6,528/1.5 bytes/min = 4,352 bytes/min. This communications rate is easily realizable through affordable, commercial wireless modems rated at 9,600 or 19,200 baud.

Figure 6.20a presents the distribution of the number of flooding messages, both sent and received, at DTMC5, as a function of the simulation time. The reasoning underlying the choice of DTMC5 is provided earlier in this chapter. For every DTMC, the maximum number of inter-DTMC flooding messages, both sent and received, is recorded and the data is graphed in Figure 6.20b. It is observed that for every DTMC, the maximum number of inter-DTMC flooding messages coincides with the beginning of simulation. The reason is as follows. As soon as vehicles are introduced on the highway segments, all of the DTMCs compute new C.M. values for the segments they control. The new C.M. values are likely to differ substantially from the default value of 100 mph, thereby requiring them to be propagated immediately to all other DTMCs. As simulation progresses, changes in the C.M. values of the segments are incremental and the frequency of flooding messages decreases. The graph in Figure 6.20b is obtained for traffic density 8, that is, 30,284 vehicles and is the result of sampling the message queue of each DTMC every simulation second. Every flooding message consists of four fields and requires 16 bytes. In Figure 6.20b, the minimum and maximum number of flooding messages are 55 and 122 respectively. The corresponding data communications rate are  $55 \times 16$  bytes/simulation second =  $55 \times 16 \div 1.5$  bytes/min = 586 bytes/min and  $122 \times 16$  bytes/simulation second =  $122 \times 16 \div 1.5$  bytes/min = 1,301 bytes/min which is also easily realizable through commercial wireless modems rated at 9,600 or 19,200 baud. Given that the resolution of the simulation is 1 simulation second or 1.5 minutes of real operation, the resolution of the data presented here is also limited to 1.5 minutes.

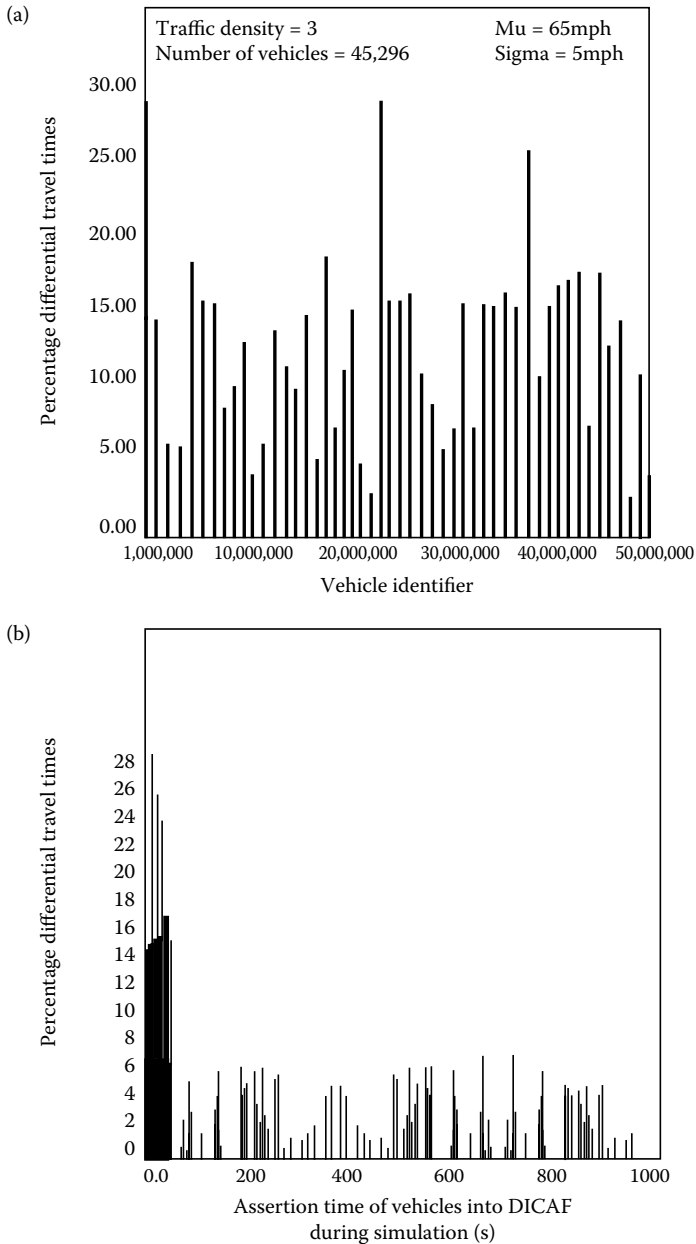
Thus, as evident from Figures 6.19 and 6.20, one of the principal objectives of DICAF namely, an affordable and cost-effective highway infrastructure, is achieved. In contrast, a centralized algorithm would theoretically require a much higher communications rate, implying expensive interfaces.

Figures 6.21a and 6.21b present results from modeling and simulating a large-scale, complex, highway system shown in Figure 6.9. The purpose of this study is to examine the applicability of the asynchronous, distributed algorithm, DICAF, to a large-scale system with 50 DTMCs and 45,296 autonomous vehicles, executing on 51 concurrently executing workstations. The graphs in Figures 6.21a and 6.21b correspond to traffic density three, represent 24 hours of actual highway operation, and traffic generator parameters of  $\mu = 65$  mph and  $\sigma = 5$  mph. Figure 6.21a yields the average and standard deviation values of 0.038% and 0.577% respectively, which



**FIGURE 6.20** (a) Number of inter-DTMC flooding messages at DTMC5 as a function of simulation time, (b) maximum inter-DTMC communication messages (send + receive) for every DTMC.

reveal that despite competition for resources among 45,296 vehicles the average vehicle’s travel time exceeds the absolute minimum by only 0.038%. Figure 6.21b reveals that, except for those that are asserted at the beginning of the simulation, all vehicles asserted into DICAf reach their destinations and that their travel times are uniform regardless of the time at which the vehicles are asserted. Other performance



**FIGURE 6.21** Modeling and simulation of the 50-DTMC highway system in DICAF. Percentage differential travel times [calculated as (actual travel time for each vehicle – ideal travel time) / (ideal travel time) × 100%] for (a) for all vehicles, (b) as a function of the assertion time of the vehicle.

graphs reveal the same general behavior as those observed for the 9-DTMC system and are not presented here.

### *Limitations of DICAf*

The issues of accidents, incidents, and the consequent congestion are not modeled in this chapter. Thus, the only source of congestion in this study arises from the stochastic speeds and destinations of the traffic that is generated and asserted into DICAf. In general, it is likely that the C.M. values for the up and down lanes for any highway segment will differ. However, DICAf assumes that the up and down lanes share a single C.M. This assumption is justified because the stochastic destinations of vehicles in DICAf will probably imply similar number of vehicles and C.M. values for both up and down lanes. This chapter is interested in uncovering general insights about the performance of DICAf under representative highway conditions, and not in accurately modeling realistic traffic in specific federal and state highways. Therefore, the traffic distribution is assumed to be independent of the time-of-day, day-of-year, and so on. As described earlier in this chapter, while every DTMCs is modeled through a workstation, each of the thousands of vehicles is modeled as a process. Processes migrate from one DTMC to a subsequent DTMC and are executed by the workstation underlying the DTMC in whose jurisdiction the corresponding vehicle is currently located. The principal reason for choosing this approach is the lack of availability of a test bed with 45,000 interconnected workstations. While this model differs from reality, it has its advantages and disadvantages. On the negative side, it does not explicitly model the procedures related to connection establishment between the DTMCs and the vehicles. Also, the exchange of data between DTMCs and vehicles in an operational system will involve actual messages that are slow relative to the exchange of local data structures in the simulation. On the positive side, the performance data presented in the chapter is conservative in that the 50 workstations bear the computational burden of every one of over 45,000 vehicles. The performance data of an actual operational system, developed based on DICAf, where each vehicle's computing engine executes its own routing functions, is very likely to be vastly superior to that predicted by the simulation reported in this chapter.

---

# 7 Stability of RYNSORD under Perturbations

## 7.1 INTRODUCTION

According to the literature, the most comprehensive treatment of stability occurs in the disciplines of physics and control systems. The motivation for defining stability is well stated by Stewart in the foreword of the 1992 translation of A.M. Lyapunov's *The General Problem of the Stability of Motion* [91]. Stewart notes that Lyapunov recognized that there are many distinct concepts of stability—different ways to formalize the idea that “small disturbances lead to small changes in the motion.” This general concept has applied to a wide range of disciplines from engineering to political science. In each case, however, the definition has been adapted to the area to which it is being applied. This chapter will take the same liberties but will remain motivated by the concept of “small disturbances lead to small changes in the motion.”

Chen [92] describes three types of stability in control theory: (i) bounded-input bounded-output stability, (ii) marginal stability, and (iii) asymptotic stability, which are presented subsequently. For each of these types, control theory allows their definitions to be expressed through differential equations, state-space, and transfer function models.

### *Bounded-Input Bounded-Output Stability*

Bounded-input bounded-output (BIBO) stability is defined as one where for every bounded input the output is also bounded. A bounded function is one whose magnitude is less than some constant for all time. The application of this definition to RYNSORD [93,94] is straightforward in that the input may be represented by the number of trains asserted into the system, at either a single station or any set of stations, at every time instance. The output consists of the time at which each train corresponding to each input leaves the system, that is, reaches its destination, which may also be viewed as a function of time. Thus, for RYNSORD to be BIBO stable, for any given bounded input function, the output function must also be bounded. This will only be true if every train asserted into the system reaches its destination in bounded amount of time—an important property of RYNSORD, which will be used subsequently in determining a steady-state input rate for the system. However, it will also be shown later that, for RYNSORD, one may always choose an input rate, which will result in an unbounded output. In general, however, the choice of BIBO-based stability is inappropriate since not all ADDM systems will necessarily lend themselves to a clear input-output relationship. Nevertheless, BIBO is an important property of RYNSORD and will be explored in greater details during steady-state analysis.



### *Marginal Stability*

Marginal stability is generally referred to as Lyapunov's definition of stability and Fuller [91] notes that this stemmed from his interest in astronomical problems. For many problems, a perturbation may not be fully dissipated but rather persists, within some bounds, for all time. For example, a particle in a circular trajectory around a point continually oscillates around it and is, therefore, marginally stable. This definition of stability is analogous to the definition of marginal stability adopted for RYNSORD.

### *Asymptotic Stability*

In contrast to marginal stability, asymptotic stability is one where the perturbation is eventually dissipated. Letov [95], in explaining Lyapunov's second theorem, notes that, under asymptotic stability, the disturbed motion converges to an undisturbed state as time progresses to infinity. Control systems engineers find this definition most appealing, and it constitutes the basis of the definition of *strong stability* in this chapter.

Casavant [96] regrets that it is difficult to apply control theory directly to distributed systems in that the mathematical methods generally used are difficult to apply to distributed systems. This stems from the complex interactions within a distributed system that defies attempts to describe their generally nonlinear behavior through a set of differential equations or transfer function unless significant simplifications are assumed. In contrast, this chapter adopts the approach that control theory is a valuable step in analyzing the properties of ADDM systems such as RYNSORD, even if the accuracy of the evaluation depends on the impact of the simplifications. Many of the basic concepts of control theory apply even where the rigorous mathematical foundations fail to apply.

In the discipline of distributed systems, the issue of stability is discussed relative to the properties of self-stabilization, correctness, that is, absence of deadlock, robustness, fault tolerance, and quality of service.

### *Self-Stabilization*

The study of self-stabilization in distributed systems was introduced by Dijkstra in 1973. He writes: "I call a system 'self-stabilizing' when, regardless of its initial state, it is guaranteed to arrive at a legitimate state in a finite number of steps" [97]. Thus, self-stabilization implies that the system is robust enough to recover from an illegal state. The definition of legitimate state is defined by Dijkstra [98] in terms of *privileges*, which are predicates based on a process' own state and that of its immediate neighbors. However, there is a difficulty in finding realistic systems that conform to this definition of legitimate states and privileges. An additional problem is that this definition is based on the identification of specific states and the identification of which global states are legitimate and which are not. This is a very difficult problem as the set of possible states can be enormous for a large and complex system.

A more general definition, without the specific definition of privileges, is given by Awerbuch et al. [99]. They write: "Self-stabilization formalizes the following intuitive goal: *despite a history of catastrophic failures, once catastrophic failures*

*stop, the system should stabilize to correct behavior without manual intervention.*” A catastrophic fault is defined as where the global state has been arbitrarily corrupted. This definition is a much more practical one yet goes well beyond Dijkstra’s definition. They only become the same if we define “correct behavior” to be a legitimate state. The definition proposed in this chapter espouse this intuitive goal but is not limited to corruptions in state. Awerbuch’s work focuses on noninteractive systems and approaches self-stabilization through periodically checking correctness and performing a re-execution whenever a fault has been found. In contrast, the RYNSORD system is one that continuously interacts with the environment and may not be stopped and re-executed due to practical considerations.

### *Stable Properties*

Related to this model of a distributed system is the definition of *stable properties* by Chandy and Lamport [100]. They note that if  $y$  is a predicate and is a stable property of a distributed system  $D$  then in a computation of  $D$ , once  $y$  is true, it remains true for the remainder of the computation. Other researchers including Venkatesan and Garg [101,102] use stable “predicate” rather than property. However, this definition of stable properties is very different in that it deals with properties defined as predicates of the system while the definition proposed here deals with stability in a system-wide perspective. For instance, deadlock is considered to be a stable property in the Reference 100. In contrast, in this chapter, if a system were to deadlock, it would result in the system being unstable. This apparent contradiction reflects the view that system stability is based on bounding the error in a system and that a deadlocked system is a case of infinite error. The assumption here is that the system is not designed to deadlock, as is the case with RYNSORD and most practical systems, so this is a characteristic of an unstable system.

### *Robustness*

Robustness is the ability to maintain correct behavior despite changes in the system. Schreiber [103] makes the distinction between robustness and fail-soft behavior by the type of errors; for robustness they are errors in the inputs and for fail-soft behavior they are faults in the system. Stankovic [104] offers a different definition, stating that “in the computer science literature, robustness normally refers to the ability of a system to handle failures.” The disagreement lies in the scope of the definition. Schreiber’s definition is limited to only errors in the input while Stankovic describes it in terms of “failures.” Meyer [105] identifies four properties for distributed real-time systems: (1) concurrency, (2) timeliness, (3) fault tolerance, and (4) degradable performance in the presence of faults. In control theory, a robust system is one that performs correctly despite perturbations in its state. This chapter is concerned with the performance impact of both failures as well as changes in input patterns. Perturbations do not necessarily imply a failure but represent any changes in the normal operating environment, and therefore, the definition proposed here encompasses robustness, fail-soft behavior, and degradable performance.

### *Fault-Tolerance*

Fault tolerance is concerned with making the system resilient against failures in the system, which is fundamentally different from the concerns of this chapter. The concern of stability is the performance after the fault, not how to recover from the fault.

### *Quality of Service*

In Reference 106, Garg et al. have defined stability for distributed applications. They have also adopted a performance perspective for stability and have chosen to use the quality of service (QoS) provided to the user as their performance index. They define a stable distributed application as one where the QoS is bounded for all time, including during the perturbation. The definition is limited in that, first, QoS attributes do not relate to RYNSORD and other systems and, second, the error during a perturbation may be unbounded.

## 7.2 FORMAL DEFINITION OF STABILITY OF RYNSORD

This section formally introduces the concept of stability for RYNSORD and presents a few definitions. The goal of this chapter is to define stability of RYNSORD in terms of a performance criteria so as to provide performance guarantees for the system in a dynamically changing environment. Ferrari [107] defines performance as an indication of how well a system, already assumed to be correct, works.

A critical concept for RYNSORD is a need for a quantitative error measurement, which is referred to as a user defined, measurable quantity. The three requirements for the error criteria are as follows: First, it is a quantifiable value. Second, conceptually, it represents the deviation of the system from some ideal, so the ideal must also be quantifiable. Third, the user desires to minimize the error quantity.

### *Definition 1*

**Error quantity:** *A quantitative measurement of the system performance which is expressed as  $error = |ideal - actual|$ . Both ideal and actual must be measurable or computable.*

The equilibrium or *steady state* for a distributed system is defined simply as the operational environment, that is, set of inputs and system resources, under which the system operates when the error is bounded by some finite constant for all time. The exact magnitude of this bound is unspecified except that it must be less than some constant that is less than infinity.

### *Definition 2*

**Steady-state:** *If a system exists in a steady state then the error of the system,  $e$ , is defined by  $e < K < \infty$  for all time where  $K$  is an arbitrary constant.*

The primary focus of defining stability is in what happens to the system, in a steady-state, following a change in the environment, termed perturbations, that are inevitable

in a real-world environment. The changes are classified into two categories: system-level perturbations and input perturbations. System-level perturbations are generally those which are considered as faults or failures. These include all forms of hardware failures as well as the arbitrary corruption of local states. Input perturbations are changes in the manner, or rate, of the input into the RYNSORD system. A perturbation is described in terms of the *assumption* it has violated. An assumption is described as a characteristic of the steady-state operating environment. Although multiple perturbations may conceivably infect the system simultaneously, this chapter limits a perturbation to a single change in the environment.

### Definition 3

**Perturbation:** *A perturbation is a violation of an assumption which is specified by the nature of the violation, the magnitude when applicable, and two time values:  $t_{pert}$  is the time at which the perturbation occurs and  $t_{pert.end}$  signifying the end of the changes to the system. Also  $t_{pert.dur} = t_{pert.end} - t_{pert}$ , is defined as the perturbation duration.*

### Definition 4

**Stability:** *If a system is in a steady-state and a perturbation occurs, it will return to a steady state as  $t \rightarrow \infty$ . Let  $K_1$  be the bounds on the original steady state and  $K_2$  be the bound on the final steady state. If  $K_2 \leq K_1$ , the system is strong stable, otherwise it is marginally stable.*

The distinction between strong and marginal stability is an important one. Given a strongly stable system, its steady state, and a perturbation, following repeated applications of the perturbation, the system will eventually return to a steady state that is either better or equal, in terms of the error bound, to its original steady state. In contrast, following a perturbation, a marginally stable system may result in a steady state with a worse error bound. Furthermore, repeated applications of the perturbation may exhibit a growing error bound. In the worst case, a periodic perturbation may either capitulate the system into instability or the error may oscillate between consecutive perturbations.

This chapter defines two related classes of stability distinguished by the perturbations to which they correspond: input stability and system-level stability.

### Definition 5

**Input stability:** *Stability related to input perturbations, that is, input rate, distribution, or magnitude.*

### Definition 6

**System-level stability:** *Stability related to system-level perturbations. The definition is inclusive of anything other than input perturbations, examples being component failures, that is, links and nodes, and component degradation, that is, dropped messages.*

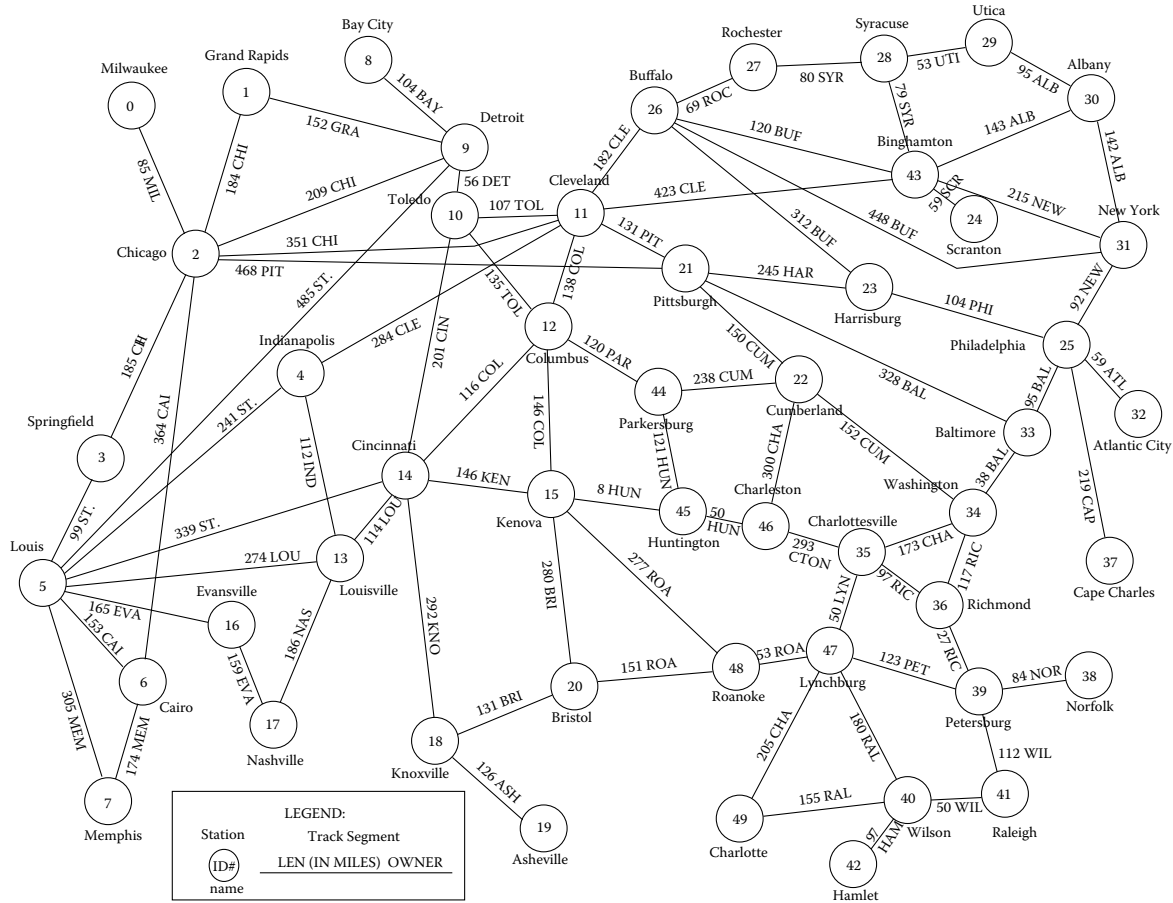


FIGURE 7.1 A representative railway network: A 50-station subset of the eastern United States railroad network.

## 7.3 MODELING RYNSORD FOR STABILITY ANALYSIS

To perform the stability analysis, RYNSORD is first modeled as a distributed discrete event simulation. This resembles a real implementation with one exception. To facilitate the simulation of a realistic system, that is, with a reasonable number of trains, while every station node is represented by a workstation, the trains are modeled as tasks and executed by the workstations underlying the stations. When a train is located at a host station, its computations are performed by the underlying workstation and its communications with other stations are carried out also through this station. When a train travels from the current station (say A) to another station (say B), the corresponding train-task in the underlying workstation for A is encapsulated through a message, propagated to B, and remanifested as a train-task in the underlying workstation at B. Thus, trains move in the simulation at electronic speeds instead of their physical speeds and a train's computation and communication subtasks are executed on the host station's underlying workstation.

While trains can propagate at speeds up to approximately 120 mph (192 km/h), the underlying, fast, computing engines of the test bed enable the simulation to execute many times faster than reality. The basic unit of time in the simulation is the timestep and it defines the finest resolution of train movements. For the reasons underlying the choice of the timestep, the reader is referred to References 93 and 94. RYNSORD permits trains to be introduced into the system asynchronously, that is, at irregular intervals of time. In addition, the trains themselves are autonomous and, therefore, their decisions are executed asynchronous with respect to each other.

### 7.3.1 IMPLEMENTATION ISSUES

For this study, a subset of the eastern United States railroad network is selected, as described in Chapter 5. A few additional tracks are added to represent a few secondary railroad segments. Figure 7.1 presents the representative railway network that consists of 50 major stations, 84 track segments, and a total length of 14,469 miles of track. A model of the network in Figure 7.1 is developed in RYNSORD with each station as a process. In addition, the lookahead for all of the experiments is set to three.

## 7.4 STABILITY ANALYSIS OF RYNSORD

### 7.4.1 ERROR CRITERIA FOR STABILITY ANALYSIS

The key performance measures in RYNSORD are the travel times required by the trains to reach their destinations and the average number of trains waiting at each station. Since stability is measured through performance behavior, this chapter proposes two error criteria, I and II, that are designed to capture the deviation of the performance measures from standard, benchmark values. This chapter proposes a novel benchmark—ideal performance measures. The ideal travel time for each train is the time required for the train to travel the shortest path from origin to destination,

in the total absence of any competing trains. The ideal number of trains waiting at any time instance at any station is zero. While these ideal values may be achieved in the total absence of any competing trains and are impractical, they are absolute minimum values and ideally suited as benchmarks. The error criteria I is expressed as

$$error = |\text{actual travel time} - \text{ideal travel time}|,$$

where the actual travel time of a train is computed as the time elapsed between the time the train is asserted into the system, in the presence of other competing trains, and the time the train reaches its destination. The error criteria II is expressed as

$$error = |\text{actual number of trains waiting at a station} - 0|.$$

Although both error criteria aim to achieve the best overall performance, they may be at odds under certain scenarios. To minimize the travel times of trains, the first error criterion may encourage a train to wait at a specific station along an optimal route until a track from it becomes available for travel. In contrast, the second error criterion may encourage the train to keep moving, while the optimal routes are occupied, through selecting longer and slower routes. Although RYNSORD selects routes based on minimizing travel time and does not directly consider the time spent waiting at the stations, both error criteria exhibit similar stability properties, reflecting the fact that the average waiting queue size and travel time for each train are related.

A real implementation of RYNSORD is a continuously running system. However, in this chapter, the simulation maintains both start and finish. Simulation is initiated with no trains in the system and the timestep set to one. The system is then executed until the end time, which for most of the experiments in this analysis correspond to 17,280 timesteps or 12 operational days. Following the initiation of simulation, the trains asserted into the system experience very little contending traffic, which appears to extend superior performance for these trains. These trains are not considered in the performance analysis of RYNSORD. Following the termination of simulation corresponding to a predetermined timestep, there may be trains still in progress in the system. These trains are marked as been asserted but never having reached their destinations, and are not considered in the performance data. In computing the results, RYNSORD only considers those trains that have successfully completed their journeys.

#### 7.4.2 STEADY-STATE ANALYSIS

Since any error analysis of a system, caused by perturbations, is relative to its normal behavior, it is imperative to first identify the steady-state behavior of the system. This section presents a steady-state analysis of RYNSORD and identifies a key criteria as the input traffic distribution. Given that freight trains dominate passenger trains in RYNSORD, this chapter assumes that the assertion of trains into RYNSORD follow a uniform distribution over time. Unlike a bursty traffic model, a uniform distribution is likely to imply a constant level of network usage, leading to efficient use of resources. At every station, the probability of a train originating at that station at each timestep,

is defined as the *input rate*. For every train originating at a station, train speeds are generated stochastically, ranging from 60 mph (96 km/h) to 100 mph (160 km/h). The final destination is also generated stochastically by assigning equal weight to every station, except the originating station, and selecting a station at random. Geographic proximity plays no part in the selection process. Since major stations, corresponding to major urban centers, are more likely to encounter high traffic densities, a set of nine “high traffic” stations are identified in Figure 7.1—Chicago, Detroit, St. Louis, Philadelphia, New York, Washington, Pittsburgh, Columbus, and Cincinnati. For the stations corresponding to these cities, the input train traffic density is set at 0.3, which, as shown later, is well above the maximum steady-state rate for the system. However, as the steady-state analysis will show, the presence of these high-traffic stations does not prevent the system from achieving a global steady state. Also, during the process of selecting final destinations of trains, these cities are assigned twice the weight of other stations to reflect that they are more likely to be selected than other cities.

A trial and error approach is utilized to determine the steady-state conditions. RYNSORD is simulated corresponding to different input rate values. Table 7.1 summarizes the average number of input trains that are generated corresponding to different input rate choices. Figures 7.2a through 7.2c present the error criterion I as a function of the assertion time, that is, the time at which a specific train is asserted into the system. In Figure 7.2a, the error does not continue to increase as time increases and, as a result, RYNSORD is considered to exhibit steady-state behavior corresponding to the input rate of 0.125. In contrast, in Figure 7.2c that corresponds to a rate of 0.175, the error clearly grows as a function of time, reflecting nonsteady-state behavior. For the input rate of 0.140, as shown in Figure 7.2b, RYNSORD exhibits both bounded behavior and growth depending on the specific stochastic input, reflecting that this input rate marks the boundary between bounded and unbounded error. As expected, different steady-state conditions exhibit different error bound values, as revealed in Figure 7.3 for the error criterion II corresponding to steady-state input rates ranging from 0.05 to 0.125.

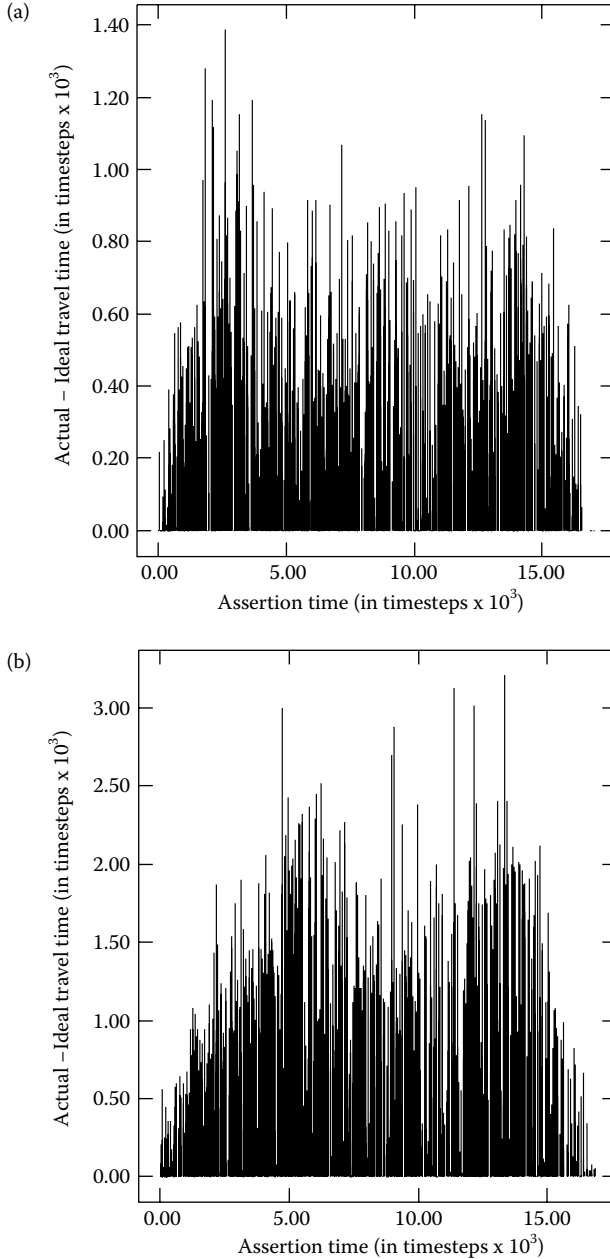
---

**TABLE 7.1**  
**Input Traffic Parameters for Steady-State Analysis of RYNSORD**

Input Traffic Density	Total Trains Introduced	Total Trains Finishing	Average Error for Completed Trains (in timesteps)	Maximum Error (in timesteps)
0.050	453	439 (97%)	67.95	439
0.100	858	822 (96%)	183.18	822
0.125	1093	1048 (96%)	309.65	1387
0.140	1250	1156 (92%)	710.25	3207
0.175	1506	1292 (86%)	1115.62	5324

---





**FIGURE 7.2** Error criterion I for each train as a function of the assertion time for (a) input rate = 0.125, (b) for input rate = 0.14, and (c) for input rate = 0.175.




FIGURE 7.2 (continued)

### 7.4.3 PERTURBATIONS TO THE INPUT RATE AND STABILITY ANALYSIS

As with any real-world system, RYNSORD is designed to execute in steady state but is likely to encounter periods of rapid fluctuation of input rates arising from any number of unforeseen circumstances. Thus, the most logical perturbation to the input rate in RYNSORD consists of an abrupt increase in the input traffic rate, sustained for a short duration. Along with the magnitude of the increase in the input rate and the length of duration of the perturbation, the choice of the steady-state operating point of RYNSORD is likely to influence the stability. It is desired that the RYNSORD design reflect a strongly stable system, that is it returns to the original steady state, at least, within finite time, following the termination of the perturbation. A number of experiments are designed and executed wherein first a steady-state RYNSORD system is exposed, one at a time, to different perturbation, under different original steady-state operating points. Second, the error criteria I and II are measured as simulation progresses and analyzed.

Table 7.2 summarizes the system characteristics under three different input rate perturbations. The magnitudes of the three perturbations are designed to push RYNSORD successively further beyond the steady-state point. Figures 7.4a through 7.4c present the error criterion I for each train as a function of its assertion time into the system, for each of the three scenarios. Figure 7.5 presents the error criterion II for each of the three scenarios as a function of the simulation time. In all of the Figures 7.4a through 7.5, the error criteria increase immediately following the perturbations. However, as

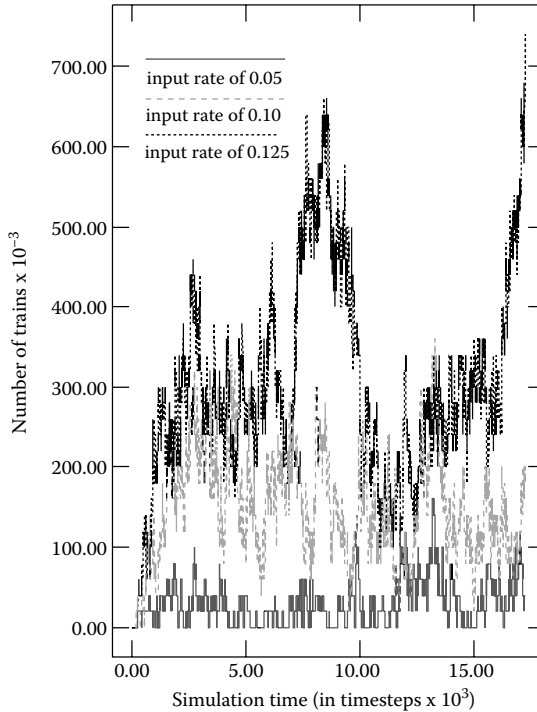


FIGURE 7.3 Error criterion II as a function of simulation time for steady-state input rates.

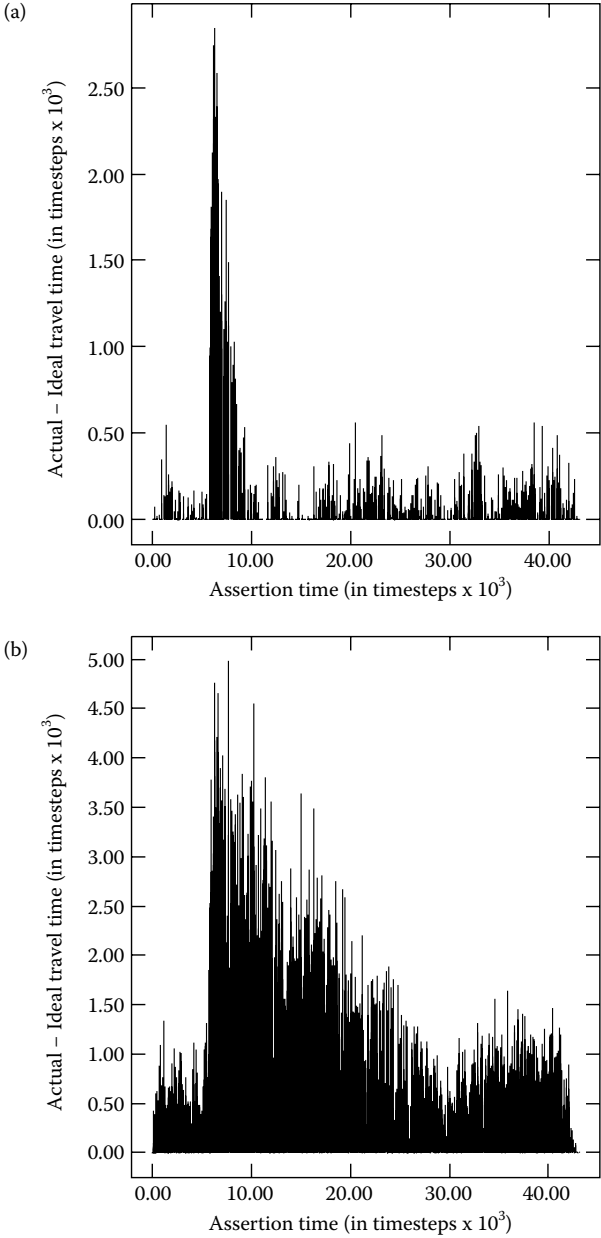
**TABLE 7.2**  
**Perturbations to Input Rate and System Characteristics**

Input Rate	Perturbation Magnitude	Perturbation Start Time (timestep)	Perturbation Duration (timestep)	Total Trains Introduced	Total Trains Finishing
0.05	+0.5	5760	720	1230	1219 (99%)
0.125	+0.5	5760	720	2955	2911 (99%)
0.125	+3.0	5760	720	3755	3443 (92%)


time progresses, the error magnitudes decrease, with RYNSORD ultimately returning to the original steady-state point for all three cases. Thus, RYNSORD is strongly stable with respect to input perturbations.

**7.4.4 PERTURBATIONS TO SYSTEM CHARACTERISTICS AND STABILITY ANALYSIS**

The basic infrastructure of RYNSORD assumes that every train is able to communicate with an appropriate station and that stations can communicate between



**FIGURE 7.4** Error criterion I for each train as a function of the assertion time for (a) input rate = 0.05 input rate and perturbation = 0.5, (b) input rate = 0.125 input rate and perturbation = 0.5, and (c) input rate = 0.125 input rate and perturbation = 3.0.



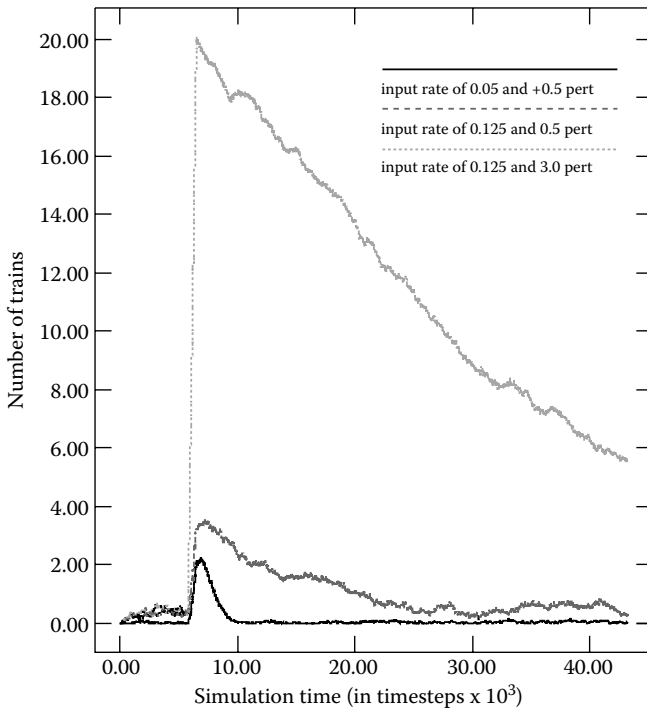
**FIGURE 7.4** (continued)

themselves. Although the RYNSORD algorithm does not explicitly take into consideration the possibility of track failures, trains are capable of determining alternate routes when one or more tracks are in use or unavailable. This section presents an investigation into the stability of the RYNSORD algorithm under such failures.

#### 7.4.4.1 Perturbations to Interstation and Train-to-Station Communications

Interactions between stations and between a train and a station constitute the key communications in RYNSORD without which trains can neither succeed in reserving tracks nor travel toward their destinations. The correctness requirement prevents a train from traveling on a track segment unless it has been granted explicit reservation. Should a reservation request, initiated by a train, remain unanswered, the train will never attempt to use the track in question. Thus, perturbations that are deliberately introduced here to affect the reservation process will bear no impact on RYNSORD's correctness.

The characteristics of the perturbations are as follows. A message propagated from one station to another never arrives at the destination. Also, a communication between a station and a train does not reach the receiver. Under such scenarios, the behavior of a train in RYNSORD is as follows. When a train does not receive a reply to its reservation request, it decides to travel on the alternate path, where available, rather than wait indefinitely for the response. When responses to both of its requests



**FIGURE 7.5** Error criterion II as a function of simulation time for different input perturbations.

for reservation fail to arrive, the train temporarily alters its lookahead to unity and renews its reservation effort. Unless both communications links along which the train propagates its reservation requests are down, which is unlikely, the most recent action by the train ensures at least one reservation response. When the unlikely event occurs, the train waits at the station and renews its reservation effort at the subsequent timestep with the restored lookahead value. Conceivably, computer communication failures are relatively short lived, and this chapter reasons that it is logical to wait for a single timestep within which the communication link is likely restored as opposed to engaging in a very roundabout detour. It is pointed out that a communication failure between two stations does not eliminate all uses of the corresponding track segment. The failure only affects those trains that attempt at reservation through the non owner station since the messages never arrive at the owner, which alone has sole capability in committing the reservation. Trains traveling from the station that owns the track are able to request and use the track.

A number of experiments are designed to measure stability: A number of different communications links are failed, different failure durations are selected ranging up to permanent failure, and different values for the input traffic rate are utilized. The objective is to analyze the impact of communications perturbation on RYNSORD and

to determine a traffic input rate for which RYNSORD is stable under perturbations. In the first experiment, two sets of three and eight links are failed separately. The links are identified subsequently through the stations at either ends. Care is exercised to avoid failing a link that is the only communication path from any station to the remainder of the network. While the set of three links are suspected high-traffic links, the choice of the set of eight links reflect the desire to distribute failures throughout the network. The simulation is executed for 17,280 timesteps for steady-state input rates of 0.05 and 0.125 respectively. The failures are asserted at timestep 5760 and last for a duration of 1440 timesteps. Thus,  $t_{pert} = 5760$ , while  $t_{pert.end} = 5760 + 1440 = 7200$ . The choice of the failure duration of 1440 timesteps, which corresponds to one full day of actual operation, reflects adequate time for repairs. In another set of experiments, the links are failed permanently, that is,  $t_{pert} = 5760$  and  $t_{pert.end} = \infty$ .

Set of three links: Baltimore(34): Washington(33); Detroit(9): Toledo(10); and Roanoke(48): Lynchburg(47).

Set of eight links: Cleveland(11): Columbus(12); Rochester(27): Syracuse(28); St. Louis(5): Detroit(9); Wilson(40): Raleigh(41); Charlottesville(35); Richmond(36); New York(31): Philadelphia(25); Knoxville(18): Bristol(20); and Parkersburg(44): Huntington(45).

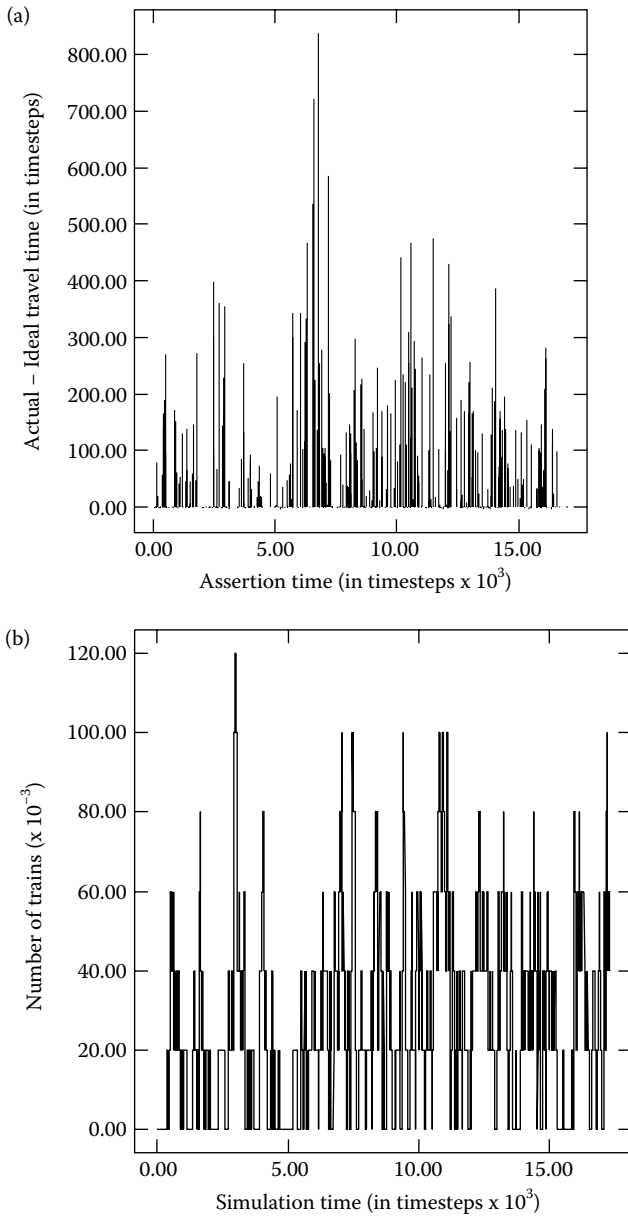
Table 7.3 summarizes the performance results and reveals that RYNSORD is strongly stable with respect to failures of finite duration. Given the higher probability of communication failures repaired quickly, the results are encouraging. However, for permanent perturbations in both sets of links, RYNSORD is observed to be marginally stable and unstable under input traffic rates of 0.05 and 0.125 respectively. Clearly, the boundary between marginal stability and instability is a function of the input traffic rate, the number of tracks failed, and the specific tracks failed. Figures 7.6a through 7.9b correspond to the set of eight failed tracks and present the error criteria I and II for different input traffic rates and perturbation durations. It is pointed out that the error criterion II mirrors the behavior of error criterion I. The results for the

---

**TABLE 7.3**  
**Performance Results for Communication Perturbations**

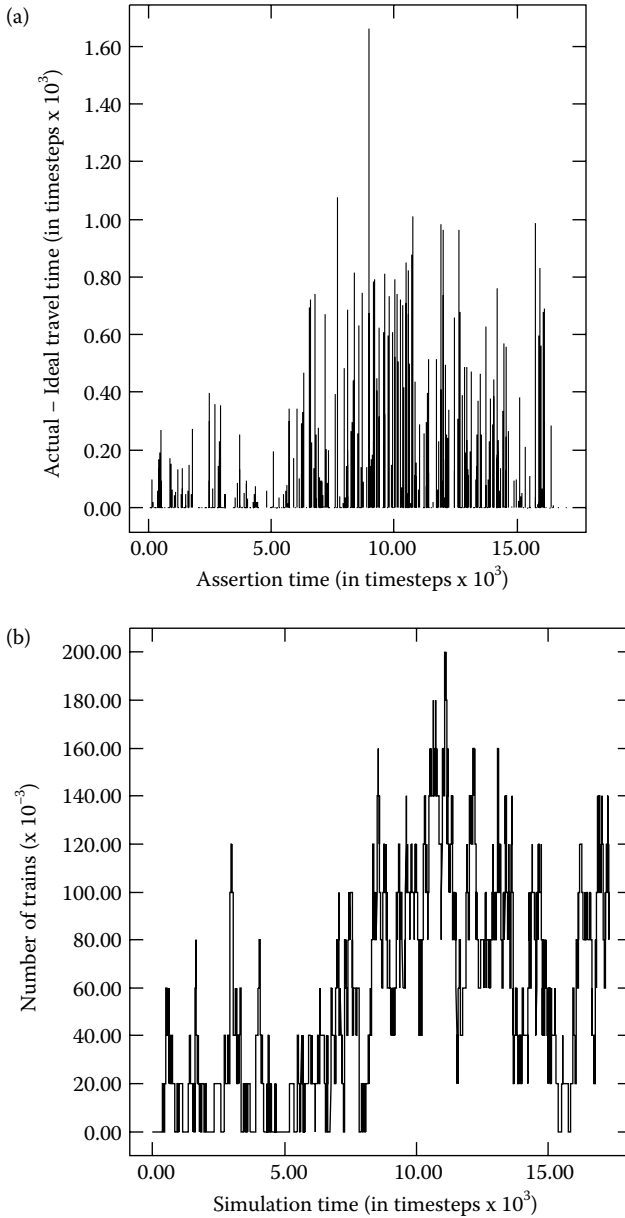
No. of Links Failed	Base Input Rate of System	Perturbation Time (in timesteps)	Perturbation Duration (in timesteps)	Stability Class
3	0.05	5760	1440	Strongly stable
3	0.125	5760	1440	Strongly stable
3	0.05	5760	$\infty$	Marginally stable
3	0.125	5760	$\infty$	Unstable
8	0.05	5760	1440	Strongly stable
8	0.125	5760	1440	Strongly stable
8	0.05	5760	$\infty$	Marginally stable
8	0.125	5760	$\infty$	Unstable

---

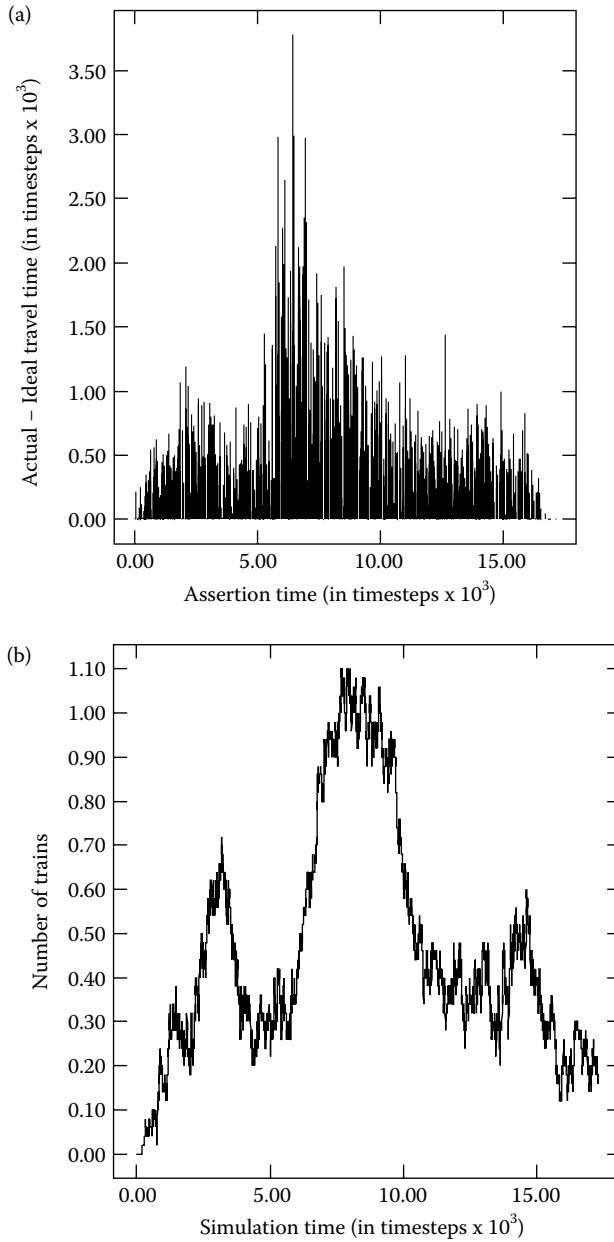


**FIGURE 7.6** (a) Error criterion I for each train as a function of the assertion time for the set of eight link failures for 1440 timesteps and input rate = 0.05, (b) error criterion II as a function of simulation time for the set of eight link failures for 1440 timesteps and input rate = 0.05.

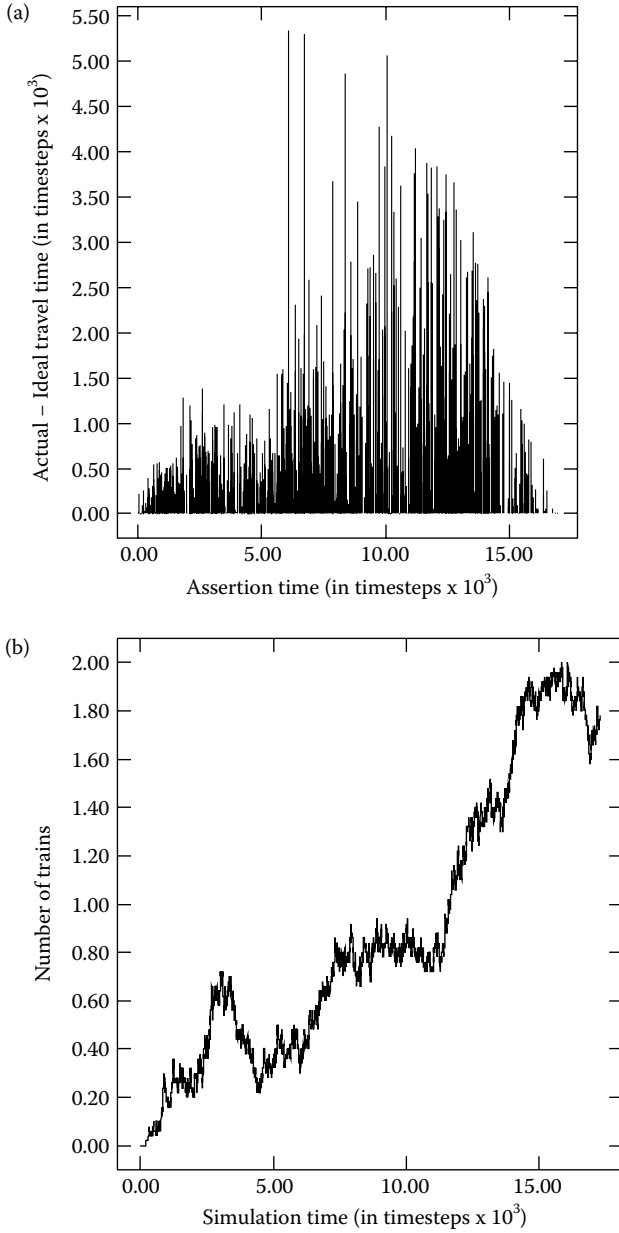




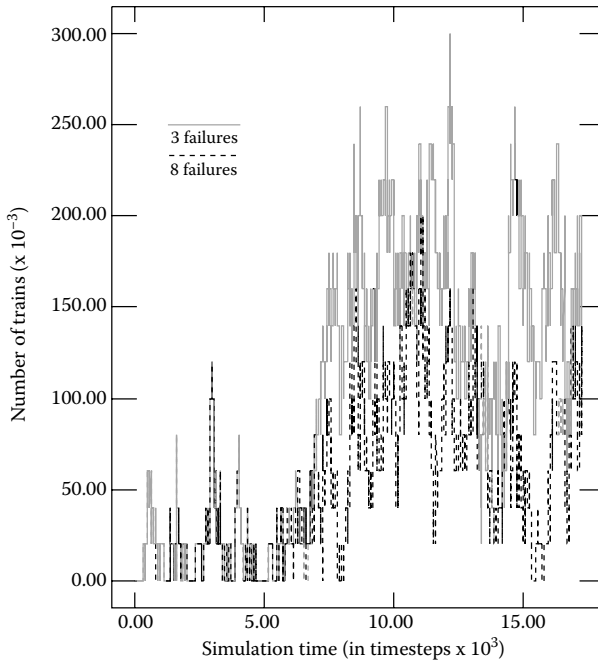
**FIGURE 7.7** (a) Error criterion I for each train as a function of the assertion time for the set of eight links failed permanently for input rate = 0.05, (b) error criterion II as a function of simulation time for the set of eight links failed permanently for input rate = 0.05.



**FIGURE 7.8** (a) Error criterion I for each train as a function of the assertion time for the set of eight link failures for 1440 timesteps and input rate = 0.125, (b) error criterion II as a function of simulation time for the set of eight link failures for 1440 timesteps and input rate = 0.125.



**FIGURE 7.9** (a) Error criterion I for each train as a function of the assertion time for the set of eight link failed permanently for input rate = 0.125, (b) error criterion II as a function of simulation time for the set of eight link failed permanently for input rate = 0.125.



**FIGURE 7.10** Error criterion II as a function of simulation time for the set of three and set of eight links permanently failed and input traffic rate = 0.05.

set of three tracks failed are similar to those for the set of eight tracks and are not presented here.

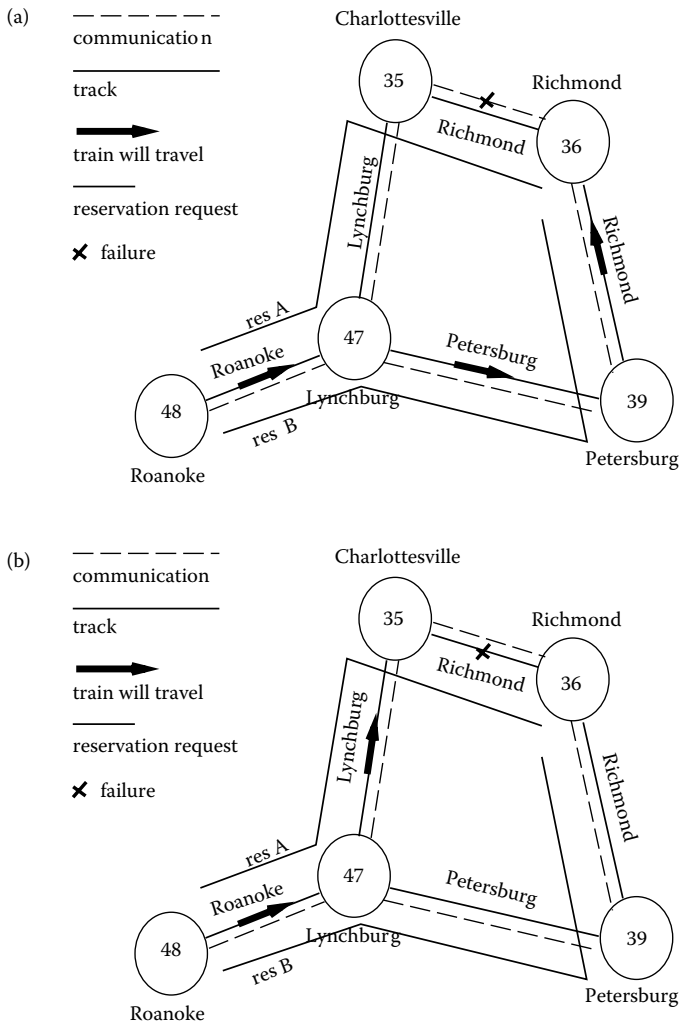
However, a comparative analysis of error criterion II for the two sets of tracks failed reveals the following, as evident through Figure 7.10. The data in Figure 7.10 corresponds to low input traffic rate and permanent failures. While RYNSORD is marginally stable for both cases, the final steady-state point for the set of three links is worse, that is, higher error bound value, relative to that for the set of eight links. The result clearly underscores the importance of the specific links failed over the number of links failed and a likely cause is the degree of congestion. Further, off-line analysis, that is, following the termination of simulation, reveals that a total of 201 trains utilized one or more of the set of three links are used in their shortest paths. In contrast, only 167 trains utilized one or more of the set of eight links. Thus, stability analysis may contribute toward identifying communications links whose failure are more likely to adversely impact the performance.

#### 7.4.4.2 Perturbations Relative to the Track Segments

A track may become unavailable following an accident, breakdown, sabotage, or due to routine maintenance. Although RYNSORD lacks elaborate mechanism to handle such failures by design, this chapter assumes the following: upon occurrence of a

failure, the stations at the two endpoints of the track segment become aware within a single timestep, that is, 60 seconds of actual operation. Also, a train already traveling on a track segment at the time of the failure will continue to travel and reach the other end safely. The stations at the endpoints will prevent future trains from using the track by canceling all reservations and by forcibly initiating reroute computations for all affected trains.

Although track failures may appear to impact the propagation of the trains similar to communication link failures, there are important differences as described subsequently. Consider Figures 7.11a and 7.11b that represent a subset of the overall net-



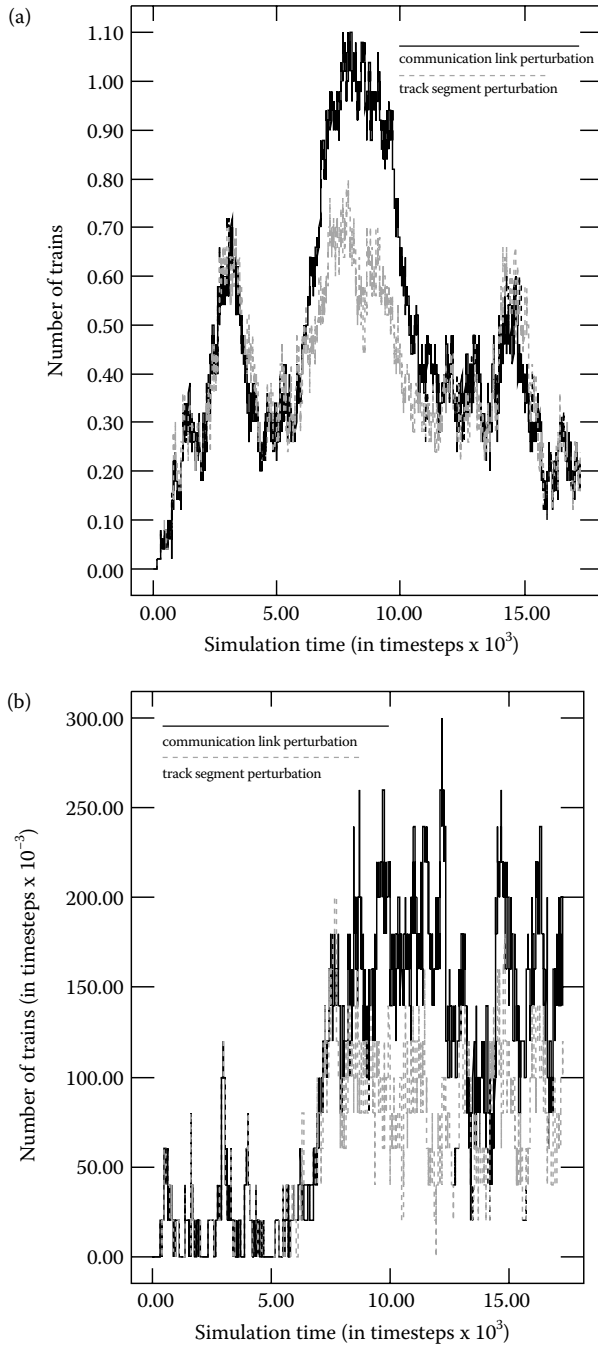
**FIGURE 7.11** (a) Illustrating communication link failure, (b) illustrating track segment failure.

work. Figure 7.11a illustrates a communication link failure between Charlottesville (35) and Richmond (36) while Figure 7.11b implies the failed track between Charlottesville (35) and Richmond (36). Consider that a train is asserted at Roanoke (48) with the destination Richmond at some time prior to the communication link failure. Reservation requests are answered while the link is good and the train chooses to follow the path A that it had requested. The propagation of the train on the track segment between Charlottesville and Richmond is guaranteed regardless of the condition of the corresponding communication link. In contrast, in Figure 7.11b, the train initiates its journey and then the track fails. The train arrives at Charlottesville and the track segment is still not repaired. The train is forced to recompute a new route and may need to backtrack to Lynchburg before continuing to Richmond via Petersburg. While this example appears to imply that track failures are likely to adversely impact the performance, subsequent analysis reveals that the impact of communication failures is more adverse.

An experiment is designed wherein the tracks corresponding to the two sets of three and eight links, described earlier in this chapter, are failed. The performance results, presented in Table 7.4, are identical to those for the communications perturbations (Table 7.3). The error-criteria graphs are also similar in behavior to those for the communications perturbations except that there are key differences. Figure 7.12a presents the error criterion II for communications link failures and corresponding track failures for the set of eight links, under high input traffic rate, and subject to a perturbation of finite duration—1440 timesteps. While RYNSORD is strongly stable relative to both kinds of perturbations and the behaviors of the graphs in Figure 7.12a are similar asymptotically, the magnitude of the error for communications perturbation is significantly worse than for track perturbation. Figure 7.12b presents the error criterion II for communications link failures and corresponding track failures for the set of eight links, under low input traffic rate, and subject to permanent perturbations. While RYNSORD is observed to be marginally stable for both scenarios and their

**TABLE 7.4**  
**Performance Results for Track Perturbations**

No. of Links Failed	Base Input Rate of System	Perturbation Time (timestep)	Perturbation Duration (in timesteps)	Stability Class
3	0.05	5760	1440	Strongly stable
3	0.125	5760	1440	Strongly stable
3	0.05	5760	$\infty$	Marginally stable
3	0.125	5760	$\infty$	Unstable
8	0.05	5760	1440	Strongly stable
8	0.125	5760	1440	Strongly stable
8	0.05	5760	$\infty$	Marginally stable
8	0.125	5760	$\infty$	Unstable



**FIGURE 7.12** Error criterion II as a function of simulation time for communications link and track perturbations for (a) high input traffic rate and finite duration perturbation, (b) low input traffic rate and permanent perturbation.

asymptotic behaviors are similar, the error magnitude corresponding to communications perturbation is considerably higher than for track perturbation. A possible explanation lies in the fact that while only the failed track becomes unavailable to a train a communication link failure may impair a train's ability to compete for reservation and therefore, travel access, for multiple tracks.

#### *Limitations of the Research*

The stability analysis of RYNSORD has revealed that it is strongly stable with respect to perturbations of finite durations to the input traffic rate and track segment failures. For permanent perturbations, the stability measure is dependent on the input traffic rate prior to the onset of the perturbation. However, it is weak with respect to communication link failures and the underlying algorithm needs redesign for superior immunity to perturbations.





---

# 8 Modeling and Simulation Techniques for ITS Designs

## 8.1 INTRODUCTION

A key characteristic of intelligent transportation system (ITS) designs, as explained in Chapter 3 and reflected in Chapters 4 through 7, is that each entity, for example locomotives, cars, and so on, carries its own computing engine while subject to transportation through the system, under asynchronous, distributed algorithm control. Every entity is viewed as an asynchronous and autonomous process with well-defined computational and communications needs. While some processes may be “stationary,” others are “mobile” within the transportation system. The exact pattern of migration of the mobile processes is dictated by the nature of the transportation system and the actual input data. The migration pattern is further complicated by the fact that every mobile process is autonomous, that is, every mobile entity determines its own migration pattern based on its unique behavior, input stimulus, and dynamic interactions with the stationary entities. Every mobile and stationary entity is characterized by unique computation and communication needs. Furthermore, the nature of the migration is asynchronous, that is, it is initiated at irregular intervals of time and may not be known a priori. Finally, in many transportation systems, the number of mobile and stationary entities is likely to be large, which, in turn, necessitates a distributed, scalable approach to modeling and simulation.

The remainder of this chapter is organized as follows: Section 8.2 introduces two competing process migration strategies that may be used to model and simulate ITS designs. Section 8.3 details their underlying software techniques while Section 8.4 presents the details of implementation of the simulation on a parallel processor test bed. Section 8.5 first presents the results obtained from executing the two simulation approaches for a representative network under realistic input conditions, followed by a comparative analysis.

## 8.2 VIRTUAL AND PHYSICAL PROCESS MIGRATION STRATEGIES FOR ITS DESIGNS

This chapter assumes the following characterization of an ITS design: (1) the number of stationary entities is relatively modest but the number of migrating entities is large, ranging from 10s to 100s; (2) the system is likely to grow in size with time requiring that the underlying approach be scalable; (3) while the stationary entities are geographically dispersed, the mobile entities are autonomous, implying that their

migration patterns are unique to every mobile process are unknown a priori; and (4) while the stationary processes are permitted to communicate directly between themselves through a static interconnection network, the mobile processes are assumed not to require direct communication between themselves for the following reasons. First, given that the number of mobile processes is large, facilities to provide direct communication between any two entities are likely to incur large overhead. This may also adversely impact scalability. Second, the underlying distributed algorithms are intelligently designed so that the stationary nodes perform the function of coordinating information between the mobile processes, when necessary. For some transportation systems, it may be necessary to provide communications between the mobile processes.

Thus, the computer model of an ITS design will consist of stationary and migrating processes executing on computing engines and mechanisms to facilitate stationary-stationary entity and mobile-stationary entity communications. Every process owns its own thread of control and is thus autonomous and asynchronous relative to other processes in the system. The capabilities of the processes are defined by the nature of the system. The stationary processes acquire necessary information from other stationary processes and mobile processes, which is subsequently downloaded and utilized by appropriate mobile processes. While the static network interconnecting the stationary processes is permanent, the mobile processes connect and disconnect dynamically and asynchronously, that is, at irregular intervals of time, with appropriate stationary processes. A migration occurs when a mobile process,  $M_i$ , chooses to disassociate itself from the stationary process,  $S_j$ , and associate itself with the stationary process,  $S_k$ , for all legitimate values of  $j$  and  $k$ .

In an operational transportation system, every stationary and mobile process is provided with its own computing engine and facilities to initiate communication with other processes. It is therefore logical to assume that in a simulation of an ITS design every stationary and mobile process will have access to its own computing engine. However, many parallel processing test beds, including the one utilized in this chapter, are likely to have far fewer available processors than the total number of stationary and mobile processes. This results in two principal strategies for representing mobile entities through processes in the test bed. They are termed virtual and physical process migration strategies and are detailed subsequently. There is a third mechanism, a variation of the physical process migration strategy, wherein, at initialization, connections are established between each mobile entity and every stationary entity. When a mobile entity needs to interact with a specific stationary unit, the corresponding connection is utilized. At other times, the connection is idle. Thus, while the overhead of dynamically establishing and destroying connections is eliminated, a maximum limit on the number of open connections per Unix process may constitute a weakness. This mechanism is not discussed in this chapter.

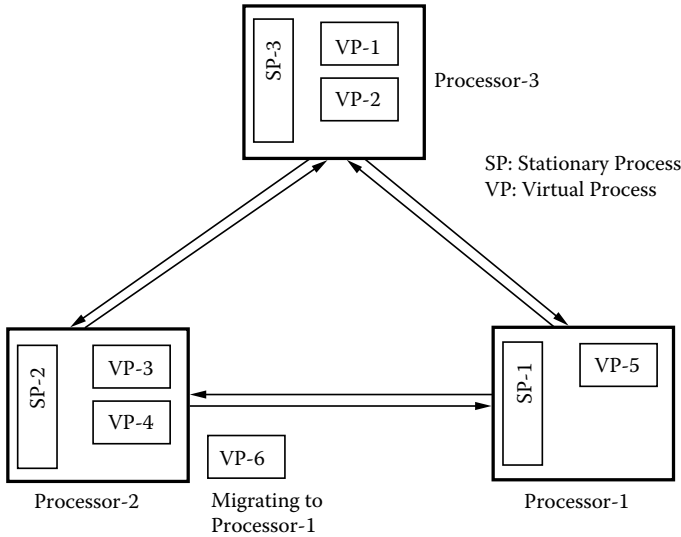
### 8.2.1 VIRTUAL PROCESS MIGRATION STRATEGY

The obvious logical choice is to represent the relatively modest number of stationary entities as actual processes, assign them to the processors of the parallel processing

test bed on an one-on-one basis, and represent the mobile entities through virtual processes. A stationary node represents an entity located at a specific geographic position. A virtual process migrates between processors, when necessary, and its computational needs are executed by the host processor underlying the stationary node, where it may happen to be located at that instant of time. By definition, a virtual process is not permanently associated with any processor. From time to time, it is associated with a processor, corresponding to a stationary node, that executes its computing needs and temporarily assigns it the status of an actual process. This strategy is termed virtual process migration (VPM) and has been discussed in Chapters 4 and 6. In Chapters 5 and 7, the VPM strategy is slightly modified in that the stationary and mobile entities are expressed through Unix processes that may be executed on a test bed with an arbitrary number of processors. Thus, unlike in Chapters 4 and 6, where the number of processors utilized equal the number of stationary entities, in Chapters 5 and 7, the number of processors required for execution need not equal the number of stationary entities. Although described earlier in the context of specific systems, this section explains VPM in detail and as a general mechanism for modeling and simulation of ITS designs. The processors are interconnected in the same topology as the stationary entities, through software protocols that are initiated at initialization time and remain unchanged throughout the simulation.

A virtual process in VPM is similar to a “thread” of an operating system. However, unlike a “thread” that contains the code, stack, stack pointer, and the program counter, a virtual process only contains the essential parameters required for its execution. The exact parameters are defined by the application program. As an example, in the modeling and simulation of the intelligent vehicle highway system (IVHS), the parameters for the mobile automobiles may include the vehicle license plate, model, manufacturer, current speed, desired speed, location, heading, origin, and destination. When a mobile entity is located at a stationary node, it “appears” at the node, that is, it is manifested as an actual process and its computing needs are executed by the host processor. Utilizing relevant information contained at the stationary node and within itself, the mobile entity determines its subsequent course of action, which may include the decision to migrate to a different stationary node. Then, the simulation migrates the corresponding virtual process with all of its parameters to the appropriate stationary node where the mobile entity again “reappears.” Thus, the behavior of a mobile entity is self-contained and is neither visible to the stationary node nor to other virtual processes that may be temporarily coresident at the same stationary node. Also, at any given time, one or more virtual processes may be resident at a stationary node and compete for the computation and communication resources. Thus, a scheduler may be utilized to assign slots of computing and communication facilities to the processes. Communication of information between the stationary process and a virtual process is achieved simply through buffer copying.

Figure 8.1 describes a simple ITS design with three stationary entities, represented through stationary processes, SP-1, SP-2, and SP-3. The simulation consists of three underlying processors—Processor-1, Processor-2, and Processor-3. At a given time instant, virtual processes VP-1 through VP-5 are resident on the three processors as shown in Figure 8.1 while VP-6 is being migrated from Processor-2



**FIGURE 8.1** Virtual process migration.

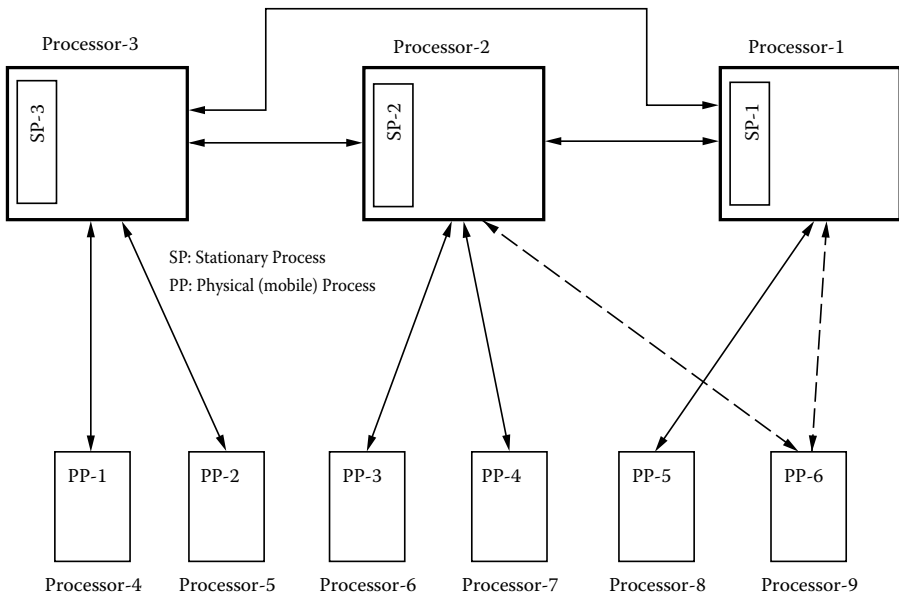
to Processor-1. During migration, the actual process corresponding to VP-6 that is resident on Processor-2 is first terminated, its essential parameters are encapsulated in a message, the message is propagated to the Processor-1, the message is decoded in Processor-2, and finally an actual process is synthesized corresponding to VP-6. Every processor is responsible for executing the stationary process and one or more virtual processes, the scheduler, and the communication primitives. A close examination yielded two important characteristics of VPM. First, in general, for a mobile entity to migrate from stationary node A to stationary node B, A and B must be connected directly. While this implies reduced complexity, it does not preclude the design of facilities to allow more complex migration. Second, the number of mobile entities at a stationary node at any given time instant is limited by the maximum number of processes permitted by the underlying operating system.

### 8.2.2 PHYSICAL PROCESS MIGRATION STRATEGY

Despite its successful use in Chapters 4 through 5, VPM incurs important limitations, which may pose difficulties in modeling future ITS systems. As the number of mobile entities increases, the competition for the host processors' computing and communication resources is likely to become acute thereby slowing down the simulation significantly. To address this limitation, this chapter proposes a competing approach, physical process migration (PPM). In PPM, every process—stationary or mobile—is allocated a unique processor. The allocation is engaged at the instant the process is initiated into the system and is disengaged when the process terminates. When a mobile process desires to communicate with a stationary process at runtime, first a

communication protocol is dynamically established between the underlying processors and then information exchange is initiated. Thereafter, when the mobile process desires to interact with a different stationary process, the old protocol is first disconnected and a new connection is established. A mobile process is allowed to maintain a connection with a single stationary process at any time. Thus, the PPM strategy is a more accurate model of reality. The static interconnection network between the stationary processes remains identical to that for the VPM. Clearly, the computational need of every mobile process is executed by its underlying processor, and where the computational needs of the mobile entities are high, there is the potential for higher efficiency and throughput relative to VPM. Unlike VPM, a mobile process may easily migrate from stationary node A to stationary node Z in PPM where a direct connection from A to Z may be lacking. PPM’s principal advantage is in the use of one processor per process. Unfortunately, this also results in a weakness in the context of the limitations of today’s test bed technology. Since test beds with 1000s of processors are not yet ubiquitous, simulation of ITS systems under PPM is limited to modest-sized mobile computing networks. PPM also inherits the limitation of high overhead for mobile-stationary process communication which includes explicit message communication following the dynamic establishment of a communications protocol.

Figure 8.2 describes the use of PPM for the simple ITS design shown earlier in Figure 8.1. In addition to the three processors that model the stationary processes, six processors, Processor-4 through Processor-9, constitute the underlying processors for the six physical processes, PP-1 through PP-6, corresponding to the six mobile



**FIGURE 8.2** Physical process migration.

entities. In Figure 8.2, solid lines between stationary and mobile processes represent protocols that are established at a given time instant while a broken line represents a protocol that is in the state of either being established or disconnected. Thus, the broken lines between VP-6 and each of Processor-2 and Processor-1 reflect the fact that the physical process PP-6 is migrating from the stationary node, SP-2, to the stationary node, SP-1.

### 8.3 SOFTWARE TECHNIQUES UNDERLYING THE PROCESS MIGRATION STRATEGIES

The static network interconnecting the stationary processes in both VPM and PPM is established during initialization of the simulation. As indicated earlier, every stationary process is assigned a distinct processor or workstation, termed node. During execution, first, a process opens a unique external input file, utilizing the identifier of the underlying workstation. This file contains the node's operating characteristics that includes its connectivity to other stationary processes. Second, the process starts to build the point-to-point connections, one at a time, utilizing the Berkeley socket protocols. When establishing a point-to-point connection between two processors, the initiator process executes a "connect" while the corresponding receiving process executes an "accept." Every connection is half duplex, implying directed edges in the network, and there may be multiple, overlapping cycles in the network. Furthermore, "connect" requires the receiving process identifier as an argument and it is non-blocking while "accept" is blocking and is designed to receive any "connect," that is, from any processor. This threatens the network initialization with the possibility of deadlock and, to counter it, the following algorithm is utilized. The underlying nodes possess unique identifiers. When a stationary process at a node (identifier X) requires connection with a stationary process at a different node (identifier Y,  $Y > X$ ), X always executes a "connect" while Y will execute an "accept." Upon completion of the network configuration in its memory, every node initiates execution of the stationary and mobile processes, which differs for the VPM and PPM strategies. The pseudocode in Figure 8.3 underscores the function at each node. In Figure 8.3, the code at label L1 first determines the set of stationary nodes in the system with identifiers higher than that of the current node. Then, the code starting at label L2 attempts to establish a connection between the current node and each of the nodes in the set through executing "connect," one at a time, until all the connections are successfully established. The code starting at label L3 corresponds to the connection establishment between the current node and all other nodes in the system with identifiers lower than that of the current node.

#### 8.3.1 SOFTWARE TECHNIQUES UNDERLYING VPM

In VPM, every mobile entity is represented through a set of parameters, which are organized into a structure. The size of the structure is a function of the application and may be dynamic. For a comparative study of the performance of VPM and

```

void make_connections()
{
L1: let t be the set of nodes with greater identifier values;
L2: while (number of elements in t != 0) {
    create socket;
    choose a node from t;
    execute a connect system call to the selected node;
    if (fails) {
        close socket;
    } else {
        send my identifier value to the selected node;
        save the socket descriptor;
        remove the node from t;
    }
}
L3: create socket;
    set up address (sockaddr_in) structure;
    bind the socket using the address structure;
    call listen relative to the socket;
    let f be the set of nodes with lower identifier values;
    while (number of elements in f != 0) {
        execute an accept system call;
        if(fails) {
            continue;
        } else {
            accept returns new socket;
            read the remote node's identifier value from the new socket;
            save new socket descriptor;
            remove the node from f;
        }
    }
    close original socket;
}

```

**FIGURE 8.3** Establishing the static interconnection network during initialization at each node in pseudocode.

PPM, the following fields are assumed for every mobile entity structure and shown in [Figure 8.4](#). The first field is the identifier of the entity, the second reflects its computational need, the third encapsulates the remaining number of messages that this entity must exchange with the host stationary process, and the fourth field stores the remaining number of hops in this entity's migration pattern.

The computational load of a mobile unit is represented by an integer, ranging from 100 to 10,000,000, and it constitutes the index of a simple “for loop.” That is, the number of iterations in the “for loop” equals the load value and the execution time of the iterative loop emulates the actual computational time. In VPM, at any given time



```

struct VPM_entity {
    int id;
    int load;
    int messages_per_hop_remain;
    int hops_remain;
};

```

**FIGURE 8.4** Structure for every mobile entity in VPM.

```

void wake_mobile_entity(struct VPM_entity)
{
L1: if (load > 0) {
    do minimum(100, load) for-loop iterations; // time slice is 100 //
    decrement load;
}
L2: if (messages_per_hop_remain > 0) {
    write to a variable location to emulate a message;
    decrement messages_per_hop_remain;
}
L3: if (load equals 0 AND messages_per_hop_remain equals 0) {
    determine stochastically where this entity must migrate;
    migrate the process and remove it from the scheduler's list;
}
}

```

**FIGURE 8.5** A mobile entity remanifested as a process at a node.

instant, one or more mobile entities may be coresident at a stationary node, competing for the single thread of control. To ensure that every mobile entity receives its fair share of the thread of control, the simulation proposes to use “time slicing,” wherein every virtual process voluntarily gives up control after executing the loop for every 100 iterations.

Upon arrival at a node, a mobile entity is remanifested as an actual process and is enqueued in the scheduler’s list. The scheduler allocates a time slot and executes the body of the mobile entity in the time slot. The pseudocode in Figure 8.5 represents the body of the mobile entity where its activity is emulated through executing iterations. The iterations are executed in sets of 100. In Figure 8.5, the statement at label L1 checks whether any iterations are left to be executed. If affirmative, a number of iterations equal to the minimum of 100 and the value of load is executed. If negative, all of the scheduled iterations have been completed. Next, the statement at label L2 checks whether any of the scheduled number of messages that need to be communicated are outstanding. If affirmative, a message communication is emulated by writing into a variable location. Otherwise, all scheduled messages have been communicated. The statement at label L3 detects the scenario that all scheduled iterations and messages have been completed. Then, the subsequent migration of the mobile entity is determined and it is propagated to the subsequent destination node.

During migration, the parameters of the mobile entity are encapsulated in a message. The address and size of the message are passed to the operating system through the “write” system call, which then writes it to the appropriate outgoing socket and executes the transfer.

At the receiving end, the node polls for the arrival of the mobile entity using the “select” system call with 0 timeout. Select maintains the ability to monitor multiple socket connections from within a single function call. When the message arrives, the node remanifests it as an actual process and enqueues it in the scheduler’s list. The pseudocode in Figure 8.6 underscores this function of the node. The code statement at label L1 in Figure 8.6 checks for a new message that encapsulates the arrival of a mobile unit. If affirmative, the message is read and the corresponding process is synthesized.

The scheduler implements round-robin scheduling of the stationary process and one or more mobile processes that may be coresident in the host processor. When it is scheduled for execution, a mobile entity is first dequeued, then executed, and then either requeued into the scheduler’s list or marked for migration to a different node. The functionality of the scheduler is shown in Figure 8.7, in pseudocode. In Figure 8.7, the statement at label L1 reflects the scheduler dequeuing the first element from the queue of mobile entities and then executes it. If the entity has exhausted its iterations and is marked for migration, as detected by the code at label L2, the scheduler encapsulates it in the form of a message and propagates it to the output. Otherwise, following execution, the entity is requeued back into the list of mobile units at the node.

The main body of the program, executed by the node, integrates all of the above functions to describe the overall operation and is shown in Figure 8.8. First, the

```
VPM_entity check_migrate()
{
    struct timeval timeout;
    timeout.tv_sec = 0
    timeout.tv_usec = 0
    fd_set fdvar;
    FD_ZERO(&fdvar);
    FD_SET(socket, &fdvar);

    L1: if (select(socket + 1, &fdvar, 0, 0, &timeout))
        VPM_entity new_entity;
        read(socket, &new_entity, sizeof(VPM_entity));
        return new_entity;
    } else {
        return 0;
    }
}
```

**FIGURE 8.6** A node intercepts a message encapsulating a mobile entity.

```

void schedule()
{
L1: dequeue the first element from the queue of waiting mobile entities;
    call wake_mobile_entities() for this mobile entity;
L2: if (mobile entity migrated) {
        write(socket, &mobile_entity, sizeof(VPM_entity));
        free the structure corresponding to the migrated mobile entity;
L3: } else
        enqueue the mobile entity back into the queue of waiting mobile entities;
}

```

**FIGURE 8.7** The function of the scheduler at each node.

```

void main()
{
L1: initialize data structures for the scheduler, etc.;
L2: synthesize mobile entities at this location;
    loop until end of simulation {
L3: call check_migrate;
        if (new mobile entity has arrived from an adjacent stationary entity) {
            assign the mobile entity to the scheduler;
        }
L4: execute schedule() to schedule the execution of the mobile entities;
}

```

**FIGURE 8.8** The main program corresponding to each node.

data structures are initialized as reflected by the statement at label L1. Then the mobile entities are synthesized utilizing information contained in the external data file, as represented by the code at label L2. The remainder of the program executes the following two operations, repeatedly and in the proper sequence: (1) Check if a new mobile node have migrated from an adjacent stationary nodes. If affirmative, remanifest the mobile entity as an actual process using the network message, and then insert the process into the scheduler's queue for execution at a later instant of time. (2) Execute the scheduler, which allocates time slices to the enqueued mobile nodes. The statements at labels L3 and L4 reflect the tasks (1) and (2) respectively.

### 8.3.2 SOFTWARE TECHNIQUES UNDERLYING PPM

#### *Mobile Entities*

At initialization, every mobile entity is associated with a processing node, which is responsible for executing the iterative loop, message transfer, and migration routines. A message transfer with respect to a stationary node requires the presence of a network protocol.

In the event of migration, first the old connection between the mobile process and a stationary process, if any, must be terminated. Normally, at either end of the

```

void migrate(stationary node)
{
    call disconnect() to terminate connection;
    determine stochastically where this process will migrate;
    try_connect(stationary node); // attempt connection to the new stationary entity /
}
void disconnect()
{
    send disconnect command to the current stationary node;
    wait for an acknowledge signal;
    call close to terminate the connection;
}
int try_connect(stationary node)
{
    call connect() to attempt connection;
    if (success) {
        send id of this mobile entity;
        return OK;
    } else {
        return ERROR;
    }
}
}

```

**FIGURE 8.9** The migration, disconnection, and reconnection functions of each mobile entity in PPM.

connection, a “close” system call is executed. However, if the execution of the system calls by the two processors are not synchronized, the connection may be closed only partially and a SIGPIPE signal is generated when a process attempts to write to it. To avoid this undesirable side effect, cooperation is required at both ends. In this chapter, when a connection is slated to be terminated, first a termination command is sent from the mobile node to the stationary node. Second, the mobile node awaits an acknowledge from the stationary node, following which both processes execute the “close” system call. Next, a new connection is established with the target stationary node. The pseudocode in Figure 8.9 describes the migration, disconnection, and reconnection functions. The migrate function executes calls to the disconnect and try\_connect functions.

The mainLoop(), shown in Figure 8.10, integrates all of the subfunctions, described earlier, to present the overall function of mobile node in PPM. First, the network topology of the stationary nodes is read from an external data file, as reflected by the statement at label L1. This information is utilized to determine migration related decisions during the simulation. Second, the internal data structures are created, represented by label L2. Third, the following three operations are executed repeatedly: (1) the iterative loop is executed to simulate actual computational operations, (2) it exchanges dummy data elements with the stationary node to which it is connected at the current time instant, and (3) it examines whether the required number of data elements have been exchanged and the loop executed for the desired number of iterations. If affirmative, the mobile entity determines a new stationary node, stochastically, and initiates migration. The statements starting at labels L3, L4, and L5 represent the tasks (1), (2), and (3), respectively.

```

void mainLoop()
{
L1: read static network topology from external data file;
L2: initialize internal data structures;
    establish connection with the default stationary node;
    loop until end of simulation {
L3:   if (load > 0) {
        do minimum(100, load) for-loop iterations;
        decrement load;
    }
L4:   if (messages_per_hop_remain > 0) {
        write(socket, default_message, length(default_message));
        decrement messages_per_hop_remain;
    }
L5:   if (load equals 0 AND messages_per_hop_remain equals 0) {
        reset the values of load and messages_per_hop_remain;
        determine a new target stationary node, stochastically.
        call migrate();
    }
    }
}

```

**FIGURE 8.10** The main program corresponding to each mobile entity in PPM.

It is pointed out that for the VPM paradigm the software pieces in Figures 8.9 and 8.10 that are executed by every mobile entity process in the PPM paradigm are contained within and executed by the stationary node. Also, the disconnect and try\_connect functions in Figure 8.9 that underlie the mobile units in PPM have no counterpart in VPM.

### *Stationary Entities*

The behavior of a stationary node includes three functions—(1) accept connection from the mobile entities, (2) exchange data with mobile entities, and (3) accept termination request from a mobile entity.

Given that a connection from a mobile node to a stationary node is initiated asynchronously and dynamically by the mobile entity, every stationary node must necessarily provide an entry point where the mobile node can initiate a connection. To realize the entry point, every stationary node binds a special socket and periodically listens to it through a select system call to determine whether a mobile node desires connection with it. The statement at label L1 in Figure 8.11 realizes this task. If affirmative, the stationary node executes an accept system call, reads and stores the identifier, and initiates the establishment of a connection with the mobile node. Upon connection, two kinds of messages are communicated—data elements and disconnection request. Data elements are exchanged between the mobile and stationary nodes and, in this chapter, the data is dummy and simply discarded. When the mobile node intends to disconnect, it propagates a disconnection message to the stationary node.

```

void mainLoop() {
    read the static network topology from the external data file;
    establish connections to other stationary nodes using make_connections();
    create and bind special socket to permit mobile nodes to connect;

    struct timeval timeout;
    timeout.tv_sec = 0
    timeout.tv_usec = 0
    loop until end of simulation {
        fd_set fdvar;
        FD_ZERO(&fdvar);
        FD_SET(mobile node connection socket, &fdvar);
L1:    if(select(mobile node connection socket + 1, 0, 0, &timeout)) {
            issue accept system call;
            read mobile node id and save new socket descriptor and id;
        }
        for(all mobile node sockets)
            FD_SET(socket, &fdvar);
        select(highest descriptor value + 1, &fdvar, 0, 0, 0)
        read message from socket;
        switch (message type) {
L2:    case disconnection:
            send acknowledge signal;
            close the socket;
            remove the socket handle from storage;
            break;
L3:    case data transfer:
            read message;
            discard message;
            break;
        }
    }
}

```

**FIGURE 8.11** The main program corresponding to each stationary node in PPM.

In turn, the stationary node will propagate an acknowledgment of the disconnection and execute the “close” system call to disconnect. The statements at labels L2 and L3 in Figure 8.11 correspond to the receipt of disconnection and data element transfer.

It is pointed out that the program executed by each stationary node in PPM, as shown in Figure 8.11, differs from that in Figure 8.8 in that it neither emulates nor executes the activities corresponding to the mobile entities. Instead, it accepts connection and disconnection requests from the mobile entities.

## 8.4 IMPLEMENTATION ISSUES

The VPM and PPM strategies are implemented on a test bed of 65+ SUN Sparc 10/40 workstations that are configured as a loosely coupled parallel processor. While each

workstation is outfitted with 32 MB of memory and executes Solaris 2.3 operating system, they are interconnected by a 10 Mbit/s Ethernet. In addition, the code design permits execution under both SUN OS 4.1.3 and the freely available Linux operating systems [108]. The code is written in C++, is approximately 2000 lines in length, and is compiled by public domain GNU g++ compiler. While the code executes in the background while user execute programs on the consoles, the data presented here is obtained from simulations that are run late at night when network load is minimal.

## 8.5 SIMULATION RESULTS AND PERFORMANCE ANALYSIS

For a comparative analysis of their performance, both VPM and PPM strategies are modeled and simulated on a parallel processing test bed. The test bed closely resembles reality and a number of experiments are designed and executed. Corresponding to an actual mobile computing network where the key parameters include the size of the static network, that is, the number of stationary nodes, the interconnection topology of the static network, the number of mobile entities, the computational load of the mobile entities, the number of messages exchanged between the mobile and stationary entities at each hop, and the migration pattern of the mobile entities, the simulation represents these parameters through independent variables. The key measure of performance is the maximum over the wall clock times required by all processors in the test bed. In the experiments, the number of entities chosen reflect the fact that the test bed is limited to 65 workstations. The number of stationary nodes range from 5 to 10, the static interconnection topology is assumed to be fully connected, and the number of mobile entities ranges from 5 to 50. The computational load of a mobile unit is represented by an integer, ranging from 100 to 10,000,000, and it constitutes the index of a simple “for loop,” as explained in Figures 8.5 and 8.10. That is, the number of iterations in the “for loop” equals the load value and the execution time of the iterative loop emulates the actual computational time. The number of data elements exchanged between a mobile and stationary entity is assumed to range from 1 to 100, where each data element is 128 bytes long dummy. The choice of the 128 byte-size reflects the message size used in Chapter 4. The message communications are also referred to in Figures 8.5 and 8.10. Every mobile entity’s migration pattern, reflected in Figures 8.5 and 8.10, is (1) stochastic, that is, randomly determined, (2) unique, that is, independent of the migration patterns of all other mobile nodes, and (3) asynchronous, that is, the mobile entity may migrate at irregular intervals of time. The only constraint imposed on the mobile units is that an entity will not immediately reconnect to the stationary node to which it was connected most recently. In the simulation, unless otherwise specified, every mobile entity connects and disconnects with the stationary nodes a total of 1000 times.

Given that the processors of a parallel processing test bed are asynchronous, their execution rates differ, and that their clocks are out of phase, the order in which events are executed in the simulation may, in general, differ from that in actual operation. To preserve the order of event execution, often, different synchronization techniques are utilized. The use of such techniques, however, constitute an artifact of the simulation

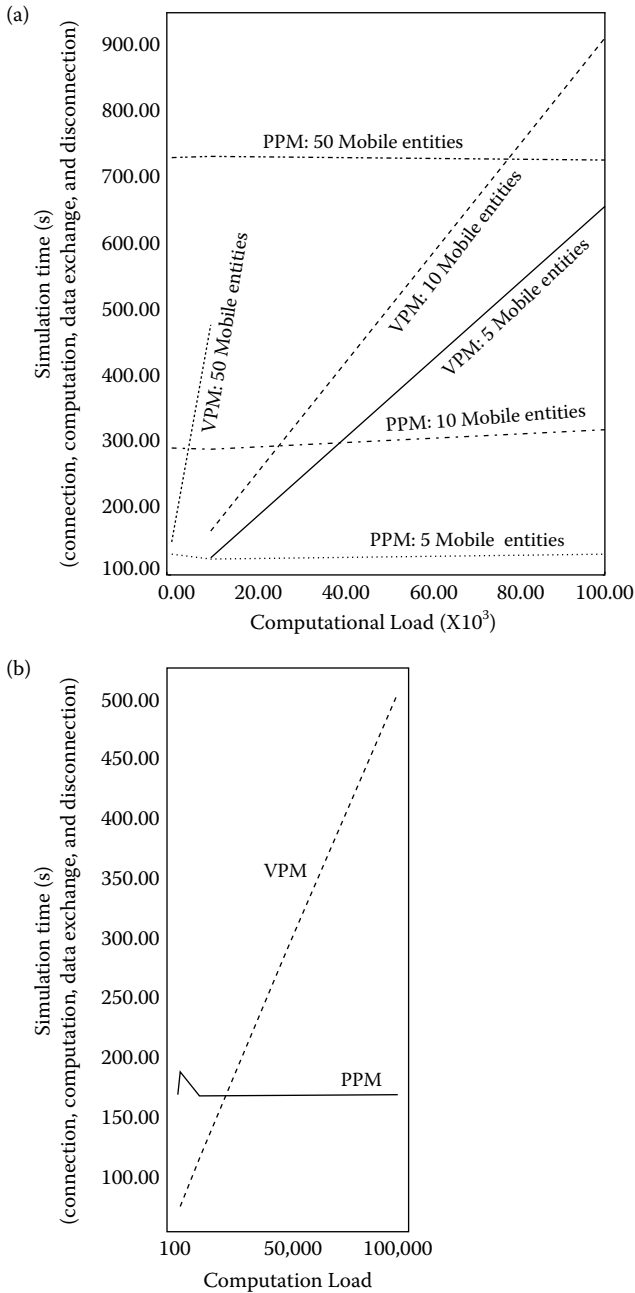
and has no correspondence in reality. These synchronization techniques slow down the simulation significantly and, as they affect both VPM and PPM similarly, the VPM and PPM implementations in this chapter are deliberately designed without them, without any loss in generality. The implementations of VPM in Chapters 5 through 7, however, utilize synchronization techniques.

The results presented here reflect a total of over 200 simulation runs, each requiring an average 1000 seconds of wall clock time, 65 concurrently executing workstations, and several MB of data collected from the simulation. When a mobile entity connects with a stationary entity, it performs computations, defined by the load value, and then exchanges data with the stationary process. The simulation terminates when every mobile entity has completed the specified number of connections and disconnections. The measured simulation time includes the time required for establishing the software protocol connection, the computation time, time for exchange of data, and disconnection time.

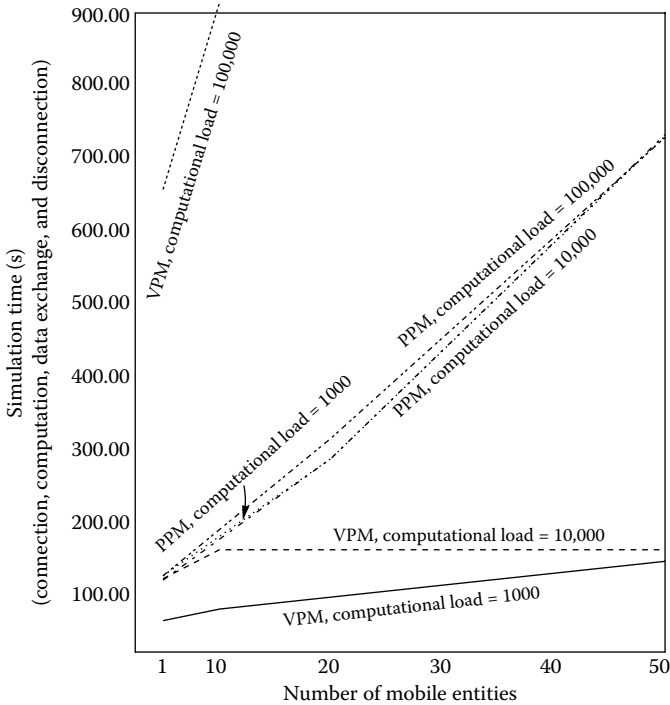
The graphs in [Figure 8.12a](#) present the variation of the simulation time as a function of the computational load of the mobile entities, for both VPM and PPM. The number of stationary nodes is set to 10 and the number of mobile nodes ranges from 5 to 10 to 50. The number of data elements exchanged at each hop is set to 1 and each mobile entity engages itself in 1000 connections and disconnections. The graphs are revealing in that while the VPM simulation times rise sharply with increasing load, the PPM simulation times remain relatively constant. For 5, 10, and 50 mobile entities, the PPM simulation times remain unchanged at 110, 300, and 710 seconds, respectively. For low computational load values, VPM exhibits superior performance due to the high overhead of connections and disconnections in PPM. For 5 mobile entities, beyond a computational load value of 20,000, PPM exhibits superior performance relative to VPM. Similarly, for 10 mobile entities, PPM's performance exceeds that of VPM for a computational load value beyond 50,000. Furthermore, while the PPM simulation executes successfully for 50 mobile entities, with load values ranging up to 100,000 and requiring 710 seconds, the VPM simulation requires extraordinarily large run times beyond load value of 10,000. The comparative behavior of PPM versus VPM is similar when the number of data elements exchanged at each hop is set to 10 for 5 stationary nodes and 5 mobile entities, as shown in [Figure 8.12b](#). Thus, for smaller number of data elements exchanged, modest number of mobile entities, and computational load under 100,000, the PPM strategy is scalable and exhibits superior performance.

[Figure 8.13](#) reorganizes the data presented in [Figure 8.12a](#) and plots the simulation time as a function of the number of mobile nodes, for different values of computational load ranging from 1,000 to 100,000. The aim is to reveal the impact of computational load on PPM and VPM performance while the number of mobile entities is varied. It may be observed that while the VPM performance for smaller computation load values of 1,000 and 10,000 exceeds the corresponding PPM graphs, the trend reverses for a high-load value of 100,000. The PPM graphs are virtually overlapping implying that the high-computational load is equitably and efficiently shared by the greater number of processors in PPM.



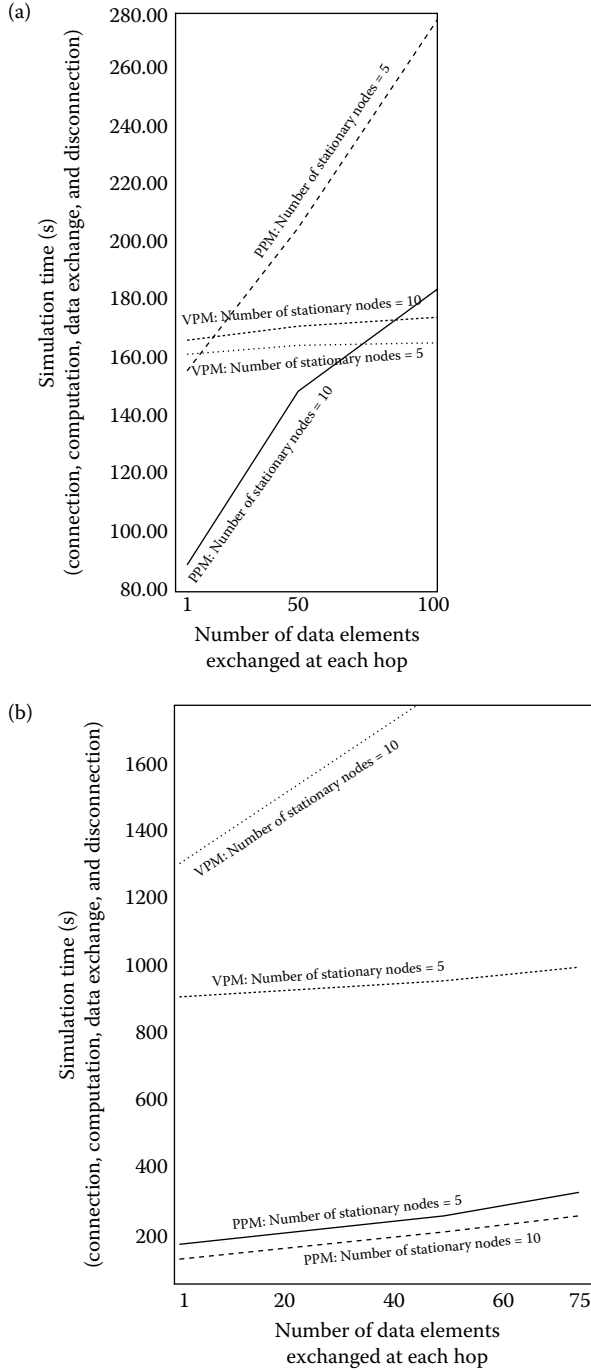


**FIGURE 8.12** VPM and PPM simulation time as a function of computational load of mobile entities, (a) 10 stationary nodes and 1 data element exchanged at each hop, (b) 5 stationary nodes and 10 data elements exchanged at each hop. The number of stationary nodes is 10, the number of data elements exchanged at each hop is 1, and the mobile unit engages in 1000 connections and disconnections.



**FIGURE 8.13** VPM and PPM simulation time as a function of the number of mobile entities for varying computational load values. The number of stationary nodes is 10, the number of data elements exchanged at each hop is 1, and the mobile unit engages in 1000 connections and disconnections.

Figures 8.14a and 8.14b plot the VPM and PPM simulation times as functions of the number of data elements exchanged at each hop. The number of mobile entities is set to 10. The computational load values are set 10,000 and 100,000 in Figures 8.14a and 8.14b respectively. The number of stationary nodes ranges from 5 to 10. The VPM graphs remain relatively uniform with increasing number of data elements exchanged since buffer copying is extremely fast. However, the observation that the VPM graph corresponding to 10 stationary nodes requires more simulation time than that for 5 stationary nodes, appears to be counterintuitive. One would have normally expected the 10 processors, corresponding to the 10 stationary nodes scenario, to finish executing the computational burden imposed by the mobile entities faster than the 5 processors corresponding to the 5 stationary nodes scenario. The reason for the observed behavior is that the function, within every stationary node, that translates the processor identifier to the socket descriptor is implemented through a linked-list. Although a general approach, the linked-list must be searched sequentially for every execution of the function. Given that the stationary nodes are fully connected, the number of sockets increase rapidly thereby slowing down the simulation for 10 stationary nodes relative to 5 stationary nodes. The Unix profiler, gprof, reveals that 40% of the simulation time is spend in the function for 10 stationary nodes as



**FIGURE 8.14** VPM and PPM simulation time as a function of number of data elements exchanged at each hop (a) load = 10,000, (b) load = 100,000. The number of mobile units is 10, and the mobile unit engages in 1000 connections and disconnections.

opposed to 23% of the simulation time for 5 stationary nodes. Since the binding between processor identifiers and socket descriptors is static, it is planned that array data structure will be utilized to provide direct and fast access.

In PPM, however, exchange of data elements involves explicit messages, and an increase in their number will require increasing simulation time, as is evident from the linear slope of the PPM graphs. For 10 stationary nodes, the PPM simulation time continues to trail the VPM simulation time up to 90 data elements exchanged per hop. Clearly, the overall computational load is executed faster by  $10 + 10 = 20$  processors in PPM relative to only 10 processors in VPM. However, when the number of data elements exchanged increases beyond 90, the overhead from explicit message passing in PPM surpasses the advantage of the greater number of computing elements and VPM supersedes PPM in performance. Where the number of stationary nodes is 5, the total number of processors at the disposal of PPM is  $5 + 10 = 15$  in contrast to 5 for VPM. Given the modest load of 10,000, the increased stress of connections and disconnections on 5 stationary processors in PPM, coupled with the high overhead of explicit message passing causes PPM to exhibit inferior performance relative to VPM. Thus, PPM loses its performance edge beyond 10 data elements exchanged at each hop.

When the computational load value is high, namely 100,000, the greater number of processors associated with PPM relative to VPM is likely to yield a superior performance for PPM. The graphs in Figure 8.14b confirm this expectation even when the number of data elements exchanged at each hop is increased from 1 to 75. Unlike Figure 8.14a where the slopes of the PPM graphs are greater than those for the VPM graphs, in Figure 8.14b, the slopes of VPM and PPM are comparable. This is due to the large computational load that lessens the influence of the message communication overhead in PPM. The PPM simulation with 10 stationary nodes employs  $10 + 10 = 20$  processors while the PPM simulation with 5 stationary nodes employs  $5 + 10 = 15$  processors, and, as expected, the performance of the former exceeds that of the latter.

In summary, both VPM and PPM play useful and effective roles in the modeling and simulation of real-world, mobile computing networks. While PPM is a more accurate model of reality, by its very nature, VPM helps realize the modeling of networks with large numbers of mobile entities on test beds with modest number of processors, although at greatly reduced performance. Chapter 6 had described an IVHS simulation with 45,000 entities representing autonomous vehicles, on a network of 60+ SUN Sparc 10 workstations. In contrast, a PPM implementation would be difficult to realize due to the large number of processors required and it would be unacceptably slow due to the large numbers of connections and disconnections. PPM is unquestionably superior when the computational loads associated with the mobile entities are high. However, its performance suffers relative to VPM where the system demands intense data exchange between the mobile and stationary units. It is hoped that with significant advances in the test bed technology in the future, with large number of processors and reduced connection and disconnection overhead, the PPM strategy may be utilized profitably.



---

# 9 Future Issues in Intelligent Transportation Systems

In our vision, the future will witness remarkable progress in intelligent transportation system (ITS) in two key fronts. First, in the near future, the key ideas underlying the intelligent transportation of matter will extend beyond vehicular traffic and trains to other modes of transportation including cargo air transport, passenger air transport, marine ferries, and personalized rapid transit (PRT) system. The distant future may even witness its extension into interplanetary travel. The need for intelligent transportation will be felt most acutely under three scenarios—increased travel speeds, significant increase in the number of travelers, and increased demand for precise and timely information by travelers, all of which are highly likely in the future. From the scientific and engineering perspective, the advances in intelligent transportation will occur in the theoretical and technological innovations. From the ordinary traveler's point of view, however, the real advance and the direct benefit will occur in the seamless and natural integration of the different modes of transportation. As a result of the integration, the traveler will (1) gain access to fairly accurate status information of any transportation mode, anywhere in the world, from any point in the system, and (2) be permitted to effect reservations, dynamically, even while en route, on any transportation mode in the world. Precision and timeliness of information are crucial to developing faith and trust in the system among the travelers, which may be delivered, in general, by distributed systems. Utilizing intelligent personalized decision aids, the traveler may process the available information to compute the most efficient route or reroute across all different transportation modes, including air, railways, automobiles, ferries, and so on. The most frequent causes for replanning include changes in the traveler's intention and needs and unscheduled delays in a currently reserved transportation system.

Second, ITS systems are complex, very expensive, and, once deployed, it is logical to expect them to remain in service for a reasonably long period of time. It is, therefore, absolutely essential to develop a sound and comprehensive understanding since such systems must be amenable to enhancements as the needs evolve with time. For this as well as for efficiency and economy, the exact details of the system architecture and design tradeoffs must be studied thoroughly, utilizing the most practical scientific tool available to us today—behavior modeling and asynchronous distributed simulation. Under asynchronous distributed simulation, a single simulation run for a highly complex system may be executed in a matter of days while the current uniprocessor simulators may require up to months. Given that a study may require up to hundreds of simulation runs, behavior modeling and asynchronous distributed simulation may

yield insights into the behaviors of such complex systems, thereby constituting an indispensable tool toward developing future ITS systems. As an example, consider the need to interconnect a number of traffic management centers in a given geographical region of the United States. A behavior modeling and simulation effort may provide meaningful and valuable insights into the topology of the interconnection and the nature of the information exchange between the centers, for a given set of long-term, high-level objectives. It would not be logical to invest millions of dollars and deploy a system, which, after completion, reveals an inability to meet the objectives and sustain growth.

The scope of research in intelligent transportation systems in the future is vast. Additional focus areas in the future range from new architecture designs for PRT and exploiting the use of embedded optical fibers in highways to assess the average speeds of vehicles to incorporating the fundamental principles of communication network design into ITS systems.

---

# 10 Description of the RYNSORD Simulator on CD-ROM and Scope of Experiments

In this chapter, we will describe the configuration and use of the RYNSORD simulator which is included in the accompanying CD-ROM. Unlike the versions described in Chapters 5 and 7, the provided software has been built for the Linux operating system. We believe that the low cost and high availability of Linux PCs make them an ideal choice of platform for distributed simulation. The system requirements include the following:

- Pentium class CPU
- 32 MB of memory
- Approximately 100 KB of disk space for the executables, while additional disk space will be required for the input and output files
- Linux glibc ELF system—kernel 2.0.36 has been tested but others should work
- TCP/IP networking enabled

In addition, the helper scripts provided assume that the rsh destination will be a trusted host so that no password will be required. Because of the security implications of this, it is not a recommended configuration for Internet-connected hosts.

If multiple hosts are used in a parallel processor configuration, the executables and input files must also be distributed either manually or via NFS.

## 10.1 INSTALLATION

1. Mount the CD as root mount/dev/cdrom/mnt/cdrom
2. Copy the rynsord tar file into your directory (e.g., /rynsord) cp/mnt/cdrom/rynsord.tar/rynsord
3. Extract it:
  - cd/rynsord
  - tar xvf rynsord.tar



#### 4. Manifest:

```

-rwxr-xr-x 1 tsl  tsl  10316 Jul 26 21:18 in_gen
drwxr-xr-x 2 tsl  tsl  1024 Jul 26 21:41 input
-rwxr-xr-x 1 tsl  tsl  4296 Jul 26 21:39 input_merge
-rw----- 1 tsl  tsl   989 Jul 26 21:55 network.10
lrwxrwxrwx 1 tsl  tsl    10 Jul 26 21:44 network.out->network.10
-rwxr-xr-x 1 tsl  tsl  19396 Jul 26 21:18 out_an
drwxr-xr-x 2 tsl  tsl   1024 Jul 26 21:56 output
-rwx----- 1 tsl  tsl   2252 Jul 26 21:53 script.10
-rwx----- 1 tsl  tsl  11466 Jul 26 21:40 script_night_4.sh
-rwxr-xr-x 1 tsl  tsl  36296 Jul 26 21:17 tsw4
-rwxr-xr-x 1 tsl  tsl  28360 Jul 26 21:17 tsy4

```

The output and input directories will be automatically created.

## 10.2 OVERVIEW

Each train station is modeled as a single Unix process (“tsw4”). The set of stations representing the simulated railway topology can be run on any collection of available host computers, from the extremes of one process per host to running all the processes on a single host. This simulation is not CPU intensive for contemporary machines and there should be no problem running on a single host. Note that the provided simulation is currently limited to 10 stations.

The typical method of operation, regardless of how many hosts are involved, is as follows (example below is for a 10-station topology, with the users home directory given as “ ” and the rynsord root as “/rynsord”):

- The appropriate network.out file is created. All processes will look for the file “/rynsord/network.out.”
- User opens 10 xterms (rsh if necessary to the different hosts)
- In each xterm, performs the following operations

```

cd ~/rynsord
setenv NETUID <uid>
setenv NETPORTNUM <portnum>

```

where uid and portnum are the unique identifier and portnumber for that node. A complete description of this is in the next section.

- The station software is started in each node:

```
tsw4 -t 10080 -l 4
```

the “-t” argument is the length of the simulation in simulated minutes, the “-l” argument is the lookahead in hops.

without these arguments, the default length is 10080, and the default lookahead is 1.

- Finally, an 11th xterm is opened

```
cd ~/rynsord
setenv NETUID 99

setenv NETUID <sync node portnum>
tsy4 -t 10080
```

This will start the synchronization node which will begin the simulation.

## 10.3 GETTING READY TO RUN

### 10.3.1 NETWORK.OUT

Mandatory file that describes the network topology. The provided example is for a 10-node network and is named “network.10.” As described earlier, this file must be copied or renamed to “network.out.” Format for this file is:

comments begin with #, everything following the # on that line is ignored.

two types of entries, node lines and link lines

node lines begin with ‘n’ and have the following syntax

```
n <uid> <name> <unused> <unused> <hostname> <portnum>
  <uid>   is the unique identifier for this node, this is an integer
          and as it implies, must be unique throughout the network
  <name>  is a string with the name of the node---this is a read
          by scanf so the standard rules of no spaces, etc, holds true.
  <unused> are unused, safest to set to ‘0’
  <hostname> Is the name of the host upon which this node will run.
             The same value returned by ‘uname -n’
  <portnum> is the portnumber upon which this node will listen for
            incoming connections. There must be a unique number for
            each individual node running on a given host. It also
            cannot conflict with any other running services. Use
            netstat -a | grep <portnum> to check for conflicts---
            the IANA reports that the following software use the
            range of 5300--5320 which is used in the provided examples.
```

link lines begin with “l” and have the following syntax

```
l <uid_A>|<uid_B> <owner> <unused> <distance>
  <uid_A> and <uid_B> are the unique identifiers of the endpoints
  <owner> is the name of the endpoint that ‘owns’ this link for
            reservation purposes
  <unused> unused, safest to set to ‘0’
  <distance> is the distance of this link in miles
```

Note that each network has the sync node as uid 99 and it is connected by links to every other node. These links are given a pseudoinfinite distance in order to prevent any routes from actually using them. Note also that the owner is not relevant for these links.

### 10.3.2 HELPER SCRIPTS

- `script.10`: is a very simple `cs`h script that is designed to open up 10 `x`terms with the environment variables set to facilitate running the simulation. `script.10` is configured to correspond with `network.10` in terms of node `uids`, names, and portnumbers. The geometry of the `x`terms is set up to work nicely on a `X`-windows setup of  $1024 \times 768$  with a virtual desktop (note that the `x` offsets start at 1040). They have been used especially for `fvwm2 @ 1024 \times 768`. Simply modify the geometry parameters as necessary.
- `script_night.10`: Is a variant of `script.10` that can be used to run the simulation in batch mode. This is very useful for doing parameter studies overnight.

### 10.3.3 INPUT GENERATION

The input files live in `/rynsord/input` and have the following format:

```
<id> <orig> <dest> <speed> <time>
  where id is the unique train identifier,
  orig is the origin stations uid,
  dest is the destination stations uid,
  speed is the train speed in mph,
  time is the time at which the train is first scheduled at the origin.
```

The program “`in_gen`” is one method of automatically generating these input files. To run it, simply type “`in_gen`” in the `/rynsord` directory. It will read the “`network.out`” file to get topology information and then prompt the user for some parameters. These include:

```
Orig vol:           The traffic volume (see Chapter 4 for more information)
Max time:           The time at which trains will cease to be asserted.
Enter spike time...: This allows the user to force a perturbation period
                    in the form of higher than normal traffic. The three
                    numbers required are time to begin, the duration, and
                    the additional volume. If no spike is desired, answer
                    ‘0 0 0’.
```

Given this information, `in_gen` will go ahead and generate the input files.

### 10.3.4 OUTPUT FILES

Each station produces its own output file called `t_output_<uid>` in `/rynsord/output`.

The format of these files follows the convention that the first letter defines the type on output line:

```
t: train information line
t [<timestamp>] id <train id> from <origin> time (<orig time>, <time left origin>)
```

Note that the final destination of each train is the node that logs the train information. For this train, the following lines will be present:

```

a <number of route computation made> <number of messages sent>
p [<A>-><B>(<time arrive B>, <time leave B>)
   this is the hop by hop path that the train took
w <time spent waiting>
   graphs! xgraph

```

At the end of each output file are the following entries:

```

1 [<A>-><B>] usage <timesteps> of <max time> per <percentage> trains <num trains>
   this is a summary of the link usage for those links owned by this node

```

For this link, the following information also applies:

```

r <reservations attempted> <reservations made> <reservations removed> <diff> <traveled>
   where <diff> is the total difference between the reserved time and the departure
   time and <traveled> is the number of times a train traversed this track segment.

```

```

R <a> <b> <c> <d> <e> <f>
   This is a summary of simulated network usage.
   <a> total number of integers sent
   <b> maximum number of integers sent in any one timestep
   <c> time at which <b> occurred
   <d> minimum number of messages sent
   <e> time at which <d> occurred
   <f> number of integers sent for waiting messages

```

Finally, at the end of the file are 10 statistics:

```

s1:   reservations originated from this node
s2:   reservations accepted at this node
s3:   reservations denied at this node
s4:   reservations processed at this node
s5:   number of reservations which are 'ideally' met
s6:   number of times the primary path was used
s7:   number of times the secondary path was used
s8:   number of times a train was allowed to leave early
s9:   total number of minutes saved by this
s10:  number of trains processed by this node

```

The utility “out\_an” is provided to do analysis of these output files. Among other things it will do is create some text files with a “.dat” extension. These are meant to be opened with the “xgraph” tool which can provide limited graphing functionality.

### 10.3.5 TROUBLESHOOTING

#### 10.3.5.1 How You Know It Is Working

There are two phases to simulation startup. The first is to complete the network initialization. A node has finished this when its output appears something like this:

```

fd 4 [0<->1] status 0
fd 5 [0<->4] status 0
fd 6 [0<->99] status 0

```

It will pause here until all the other nodes have also completed network initialization.

The second phase then begins, which is the actual simulation. The only obvious output will be the sync node, which will be giving constant “TIME UPDATE” messages as simulation time progresses. The other nodes will give period train messages as they process trains and messages.

### 10.3.5.2 Problems

If not all of the nodes complete network initialization properly, then it is mostly an error in either the network.out file or the environment variables used. Double check each xterm with the contents of the network.out file and also verify that the same portnumber is not used twice on the same machine.

If everything looks good, then it is possible that there is contention for a port. The command “netstat -a” can be used (in conjunction with grep) to check for conflicts with the portnumbers that you have selected.

If you kill (control-C or via “kill”) a simulation (and you must kill all nodes before restarting), there may be a time delay before you can reuse the same portnumbers. The “netstat -a” command will show these as being in TIME\_WAIT state. They will become available again after a short time.

If the simulation appears to freeze after running for sometime, then one possibility is that the network graph is not fully connected.

### 10.3.6 TRACK AND COMMUNICATION FAILURES

The provided simulator has the capability of simulating a single track or communication link failure. To initiate this, simply add the following command line arguments to all nodes (except the sync node):

```
-k <A> <B> <time to initiate> <duration>
    to fail a track segment between uid A and B.

-f <A> <B> <time to initiate> <duration>
    to fail a communication link between uid A and B.
```

Note that you cannot use both flags concurrently.

## 10.4 CONCLUSIONS

The information provided in this chapter, along with the accompanying software, should be sufficient for the reader to perform a set of simple experiments with the RYNSORD algorithm. Note that this is research software and may not be completely robust. This is definitely not supported software. However, if you do experience problems or have comments, feel free to mail us at “sumit.ghosh@ieee.org.” We will try our best to get back to you in a timely fashion.

---

# Bibliography

1. S. Yamanouchi. Essential information systems for railways and intensive application of ADS technology—COSMOS and ATOS. In *Keynote Address: International Symposium on Autonomous Decentralized Systems*, pp. 2–9, Tokyo, October 1999.
2. W.A. Carley. Railroads test satellite positioning in effort to improve safety, efficiency. *The Wall Street Journal Interactive Edition*, June 29, 1998.
3. T. Lomax. Five most congested urban areas: LA, DC, SFO-Oakland, Miami, Chicago. Texas Transportation Institute, June 24, 2002. Vol. 34, No. 1, <http://mobility.tamu.edu>.
4. History Channel. The Crumbling of America. June 29, 2009.
5. S. Ghosh and T. Lee. *Principles of Modeling and Asynchronous Distributed Simulation of Complex Systems*. IEEE Press: Picataway, NJ, 2000.
6. P. Heck and S. Ghosh. A study of synthetic creativity through behavior modeling and simulation of an ant colony. *IEEE Intelligent Systems*, 15(6):58–66, November/December 2000.
7. P. Heck and S. Ghosh. The design and role of synthetic creative traits in artificial ant colonies. *Journal of Intelligent and Robotic Systems*, 33(4):343–370, April 2002.
8. S. Ghosh. Creative thinking. *IEEE Circuits and Devices*, 20(1):52–62, January/February 2004.
9. D.T. Emerson. The work of Jagadis Chandra Bose: 100 years of millimeter-wave research. *IEEE Transactions on Microwave Theory and Techniques*, 45(12):2267–2273, December 1997.
10. D.E. Barrett. Traffic-noise impact study for least Bell’s vireo habitat along California State Route 83. *Transportation Research Record*, 1559:3–7, 1996.
11. D.E. Barrett. California study addresses issues of roadway noise control for the benefit of endangered songbirds. *Road Management and Engineering Journal, TransSafety*, <http://www.usroads.com/journals>, pp. 1–6, October 1997.
12. The Learning Channel. Kansai International Airport. November 26, 2000, <http://www.tlc.com>.
13. Xcel Energy Colorado System. *PHEV Impacts on the Xcel Energy Colorado System: Executive Summary*. Colorado, December 31, 2008.
14. T. Markel and K. Bennion. Xcel PHEV/V2G Project. In National Renewable Energy Laboratory, October 24, 2008.
15. W.L. Grossman. History of transportation. In *Colliers Encyclopedia*, CD-ROM. February 28, 1996.
16. S. Simon (Interviewer). Big dig—Most expensive road in U.S. history [transcript]. National Public Radio, Weekend Edition, June 22, 1996.
17. B. Blide. Personal rapid transit study in Gothenburg, Sweden. *Transportation Research Record*, 1451:87–94, December 1994.
18. I. Andreasson. Vehicle distribution in large personal rapid transit systems. *Transportation Research Record*, 1451:95–99, December 1994.
19. E.J. Anderson. Dependability as a measure of on-time performance of personal rapid transit systems. *Journal of Advanced Transportation*, 26(3):201–212, Winter 1992.
20. Raytheon Electronic Systems. Personalized Rapid Transport, PRT 2000. Technical report, Marlborough, MA, October 1996.

21. D.C. Coll, A.U. Sheikh, R.G. Ayers, and J.H. Bailey. The communications system architecture of the North American Advanced Train Control System. *IEEE Transactions on Vehicular Technology*, 39(3):244–255, August 1990.
22. N. Utamaphethai and S. Ghosh. DICAF, A high performance, distributed, scalable architecture for IVHS utilizing a continuous function—“congestion measure”. *IEEE Computer*, 31(3):78–84, March 1998.
23. Federal Aviation Authority. <http://www.faa.gov/freeflight/>.
24. Intelligent Transportation Society of America and US Department of Transportation. *ITS Architecture Development Program Phase I*. ITS America: Washington, DC, 1994.
25. S. Ghosh and T. Lee. *Intelligent Transportation Systems: New Principles and Architectures*. CRC Press: Boca Raton, FL, January 2000, Reprinted June 2002.
26. P. Varaiya and S. Shladover. Sketch of an IVHS architecture. *PATH Research Report*, UCB-ITS-PRR-91-03, February 1991.
27. S.M. Graff and P. Shenkin. A computer simulation of a multiple track rail network. In *Sixth International Conference on Mathematical Modeling*, St. Louis, MO, August 4–7, 1987.
28. K. Fukumori, H. Sano, T. Hasegawa, and T. Sakai. Fundamental algorithm for train scheduling based on artificial intelligence. *Systems and Computers in Japan*, 18(3): 52–63, 1987.
29. V.P. Rao and P.A. Venkatachalam. Microprocessor-based railway interlocking control with low accident probability. *IEEE Transactions on Vehicular Technology*, 35(3): 141–147, August 1987.
30. S.K. Kashyap. Telecommunications on Indian Railways. *IETE Technical Review (India)*, 5(3):117–119, March 1988.
31. G. Vernazza and R. Zunino. A distributed intelligence methodology for railway traffic control. *IEEE Transactions on Vehicular Technology*, 39(3):263–270, August 1990.
32. J.T. Rayfield and H.F. Silverman. Operating system and applications of the Armstrong multiprocessor. *IEEE Computer*, 21(6):38–52, 1988.
33. Interstate Commerce Commission. Ninety-ninth annual report transport statistics in the United States, December 1987.
34. P.H. Bernard. ASTREE: A global command, control, and communication system. *Transportation Research Record*, Transportation Research Board: Washington, DC 20418: 133–139, 1991.
35. R.J. Hill, S.L. Yu, and N.J. Dunn. Rail transit chopper traction interference modeling using the spice circuit simulation package. *IEEE Transactions on Vehicular Technology*, 38(4):237–246, 1989.
36. R.G. Ayers. Selection of a forward error correcting code for the data communication radio link of the advanced train control-system. *IEEE Transactions on Vehicular Technology*, 38(4):247–254, 1989.
37. A.U.H. Sheikh, D.C. Coll, R.G. Ayers, and J.H. Bailey. ATCS—Advanced train control-system radio data link design considerations. *IEEE Transactions on Vehicular Technology*, 39(3):256–262, 1990.
38. R.J. Hill. Optimal construction of synchronizable coding for railway track circuit data-transmission. *IEEE Transactions on Vehicular Technology*, 39(4):390–399, 1990.
39. Y.R. Shayan, L.N. Tho, and V.K. Bhargava. Design of Reed-Solomon (16,12) codec for North-American advanced train control-system. *IEEE Transactions on Vehicular Technology*, 39(4):400–409, 1990.

40. S. Ghosh and M-L. Yu. An asynchronous distributed approach for the simulation and verification of behavior-level models on parallel processors. *IEEE Transactions on Parallel and Distributed Systems*, 6(6):639–652, June 1995.
41. J. Misra. Distributed discrete-event simulation. *Computing Surveys*, 18(1):39–65, March 1986.
42. J.L. Hennessy and D.A. Patterson. *Computer Architecture: A Quantitative Approach*. Morgan Kaufmann: Palo Alto, CA, 1990.
43. W.R. Stevens. *Unix Network Programming*. Prentice Hall: New Jersey, 1993.
44. E.W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1:269–271, 1959.
45. I. Catling and B. McQueen. Road transport informatics in Europe—Major programs and demonstrations. *IEEE Transactions on Vehicular Technology*, 40(1):132–140, February 1991.
46. H. Kawashima. Two major programs and demonstrations in Japan. *IEEE Transactions on Vehicular Technology*, 40(1):141–146, February 1991.
47. G. King. Driver performance in highway navigation tasks. *Transportation Research Record (TRB)*, National Research Council, Washington, DC, 1093:1–11, 1986.
48. Rhode Island Department of Transportation. *State of Rhode Island Incident Management Plan*. State Office Building: Providence, RI, September 1992.
49. C. Levesque. Rhode to the future. *Traffic Technology International*, 38–42, June/July 1998.
50. J. Peters, M. McGurrin, D. Shank, and M. Cheslow. Estimate of transportation cost savings from using intelligent transportation system (ITS) infrastructure. *ITE Journal*, 67(11):42–47, November 1997.
51. D.A. Haver and P.J. Tarnoff. Future-directions for traffic management systems. *IEEE Transactions on Vehicular Technology*, 40(1):4–10, 1991.
52. R.E. Fenton and R.J. Mayhan. Automated highway studies at the Ohio State University—An overview. *IEEE Transactions on Vehicular Technology*, 40(1):100–113, 1991.
53. W.B. Powell. Optimization models and algorithms—An emerging technology for the motor carrier industry. *IEEE Transactions on Vehicular Technology*, 40(1):68–80, 1991.
54. T.M. Batz. The utilization of real-time traffic information by the trucking industry. *IEEE Transactions on Vehicular Technology*, 40(1):64–67, 1991.
55. D.H. Roper and G. Endo. Advanced traffic management in California. *IEEE Transactions on Vehicular Technology*, 40(1):152–158, 1991.
56. W. Kremer, D. Hubner, S. Hoff, T.P. Benz, and W. Schafer. Computer-aided design and evaluation of mobile radio local area networks in RTI/IVHS environments. *IEEE Journal of Selected Areas in Communications*, 11(3):406–421, April 1993.
57. S. Sakagami, A. Aoyama, K. Kuboi, S. Shirota, and A. Akeyama. Vehicle position estimates by multibeam antennas in multipath environments. *IEEE Transactions on Vehicular Technology*, 41(1):63–67, February 1992.
58. T.M. Hussain, T.N. Saadawi, and S.A. Ahmed. Overhead infrared sensor for monitoring vehicular traffic. *IEEE Transactions on Vehicular Technology*, 42(4):477–483, November 1993.
59. J.L. Kim, J-C.S. Liu, P.I. Swarnam, and T. Urbanik. The areawide real-time traffic control (ARTC) system: A new traffic control concept. *IEEE Transactions on Vehicular Technology*, 42(2):212–224, May 1993.
60. R. Von Tomkewitsch. Dynamic route guidance and interactive transport management with ALI-SCOUT. *IEEE Transactions on Vehicular Technology*, 40(1):45–50, 1991.



61. R.W. Denney and M.J. Chase. True distributed processing in modular traffic signal systems—San Antonio Downtown System. *Transportation Research Record*, 1324: 130–136, 1991.
62. D.W. Kline and P. Fuchs. The visibility of symbolic highway signs can be increased among drivers of all ages. *Human Factors*, 35(1):25–34, 1993.
63. D.I. Robertson and R.D. Bretherton. Optimizing networks of traffic signals in real-time—The SCOOT method. *IEEE Transactions on Vehicular Technology*, 40(1):11–15, 1991.
64. W.C. Collier and R.J. Weiland. Smart cars, smart highways. *IEEE Spectrum*, pp. 27–33, April 1994.
65. S.E. Shladover, C.A. Desoer, J.K. Hedrick, M. Tomizuka, J. Walrand, W.B. Zhang, D.H. McMahon, P. Huei, S. Sheikholeslam, and N. Mckeown. Automatic vehicle control developments in the PATH program. *IEEE Transactions on Vehicular Technology*, 40(1):114–130, 1991.
66. W.H. von Aulock. Smart cars may not be smart. *IEEE Spectrum*, pp. 17–18, March 1994.
67. H.D. Goldstein. *Field Test Report*. NYNEX Assurance Services: Elmsford, NY, May 1994.
68. H.J.M. VanGrol and A.F. Bakker. Special-purpose parallel computer for traffic simulation. *Transportation Research Record*, 1306:40–48, 1991.
69. D.J. Dailey, M.P. Haselkorn, and D. Meyers. Structured approach to developing real-time, distributed network applications for ITS deployment. *ITS Journal*, 3(3):163–180, 1996.
70. N. Jing, Y-W. Huang, and E.A. Rundensteiner. Hierarchical encoded path views for path query processing: An optimal model and its performance evaluation. *IEEE Transactions on Knowledge and Data Engineering*, 10(3):409–432, May–June 1998.
71. Z.A. Talib, N.S. Love, J.C. Gealow, G. Hall, L. Masaki, and C.G. Sodini. Massively parallel image processing system for intelligent transportation system application. In *Proceedings of the 1997 IEEE Conference on Intelligent Transportation Systems*, pp. 367–372, November 9–7, 1997.
72. A. Ziliaskopoulos, D. Kotzinos, and H.S. Mahmassani. Design and implementation of parallel time-dependent least time path algorithms for Intelligent Transportation Systems applications. *Transportation Research Part C: Emerging Technologies*, 5(2):95–107, April 1997.
73. N. Roupail, S.R. Ranjithan, W. El-Dessouki, T. Smith, and D.E. Brill. Decision support system for dynamic pre-trip route planning. In *Proceedings of the 1995 4th ASCE International Conference on Applications of Advanced Technologies in Transportation Engineering*, pp. 325–329, Capri, Italy, June 27–30, 1995.
74. The Minnesota Department of Transportation. The Minnesota Department of Transportation's Traffic Management Center (TMC). Minneapolis, MN, January 1994.
75. J. Upchurch, D. Powell, and P. Pretorius. Model deployment of Intelligent Transportation Systems: The AZTech experience. In *Proceedings of the 1998 5th ASCE International Conference on Applications of Advanced Technologies in Transportation*, pp. 367–372, Newport Beach, CA, April 26–28, 1998.
76. G.L. Chang, T. Junchaya, and A.J. Santiago. A real-time network traffic simulation model for ATMS applications: Part II—Massively parallel model. *IVHS Journal*, 1(3):243–259, 1994.
77. T. Junchaya and G.L. Chang. Exploring real-time traffic simulation with massively parallel computing architecture. *Transportation Research, Part C: Emerging Technologies*, 1(1):57–76, March 1993.

78. R.W. Hall. Route choice and advanced traveler information systems on a capacitated and dynamic network. *Transportation Research Part C: Emerging Technologies*, 4(5): 289–306, 1996.
79. Washington State Department of Transportation. *Seattle Area Traffic Conditions*. <http://www.wsdot.wa.gov/regions/northwest/NWFLOW/> (accessed 1998).
80. J. Clark and G. Daigle. Importance of simulation techniques in its research and analysis. In *Proceedings of the 1997 Winter Simulation Conference*, pp. 1236–1243, Atlanta, GA, Dec 7–10, 1997.
81. L. Bodin, B. Golden, A. Assad, and M. Ball. Vehicle routing of vehicles and crews: The state of the art. *Computers and Operations Research*, 10(2): 1983.
82. R.W. Allen, D. Ziedman, T.J. Rosenthal, A.C. Stein, J.F. Torres, and A. Halati. Laboratory assessment of driver route diversion in response to in-vehicle navigation and motorist information systems. *Transportation Research Record*, 1306:82–91, 1991.
83. P. Davies and G. Klein. Field trials and evaluations of radio data system traffic message channel. *Transportation Research Record*, 1324:1–7, 1991.
84. B. Kamali. Wireless communications: The nervous system of intelligent transportation systems. *Applied Microwave and Wireless*, 9(5):6–6, September–October 1997.
85. H. Sodeikat. Universal system architecture. In *Proceedings of the Pacific Rim Transtech Conference*, pp. 251–257, Washington State Conventional Center, July 25–28, 1993.
86. E. Kreyszig. *Advanced Engineering Mathematics*. John Wiley and Sons: New York, 1967.
87. A. Bonde and S. Ghosh. A comparative study of fuzzy versus “fixed” thresholds for a robust queue management in cell-switching networks. *IEEE/ACM Transactions on Networking*, 2(4):337–344, August 1994.
88. D. Bertsekas and R. Gallager. *Data Networks*. Prentice Hall: Upper Saddle River, NJ, 1992.
89. Rhode Island Department of Transportation. *State Highway Map of Rhode Island—1993 Annual 24 Hour Average Daily Traffic*. State Office Building: Providence, RI, 1994.
90. T. Lee, S. Ghosh, J. Lu, X. Ge, A. Nerode, and W. Kohn. A mathematical framework for asynchronous, decentralized, decision-making algorithm with semi-autonomous entities: Synthesis, simulation, and evaluation. *IEICE Transactions on Fundamentals*, E-83A(7): 1381–1395, July 2000.
91. A.M. Lyapunov. *The General Problem of the Stability of Motion*. Taylor & Francis: London, 1992.
92. C-T. Chen. *Analog and Digital Control System Design: Transfer-Function, State-Space, and Algebraic Methods*. Saunders College: New York, 1993.
93. T. Lee and S. Ghosh. RYNSORD: A novel, decentralized algorithm for railway networks with “soft reservation.” In *Proceedings of the 48th IEEE Annual Vehicular Technology Conference (VTC’98)*, May 18–21 1998.
94. T. Lee and S. Ghosh. RYNSORD: A novel, decentralized algorithm for railway networks with “soft reservation.” *IEEE Transactions on Vehicular Technology*, 47(4): 201–222, November 1998.
95. A.M. Letov. *Stability in Nonlinear Control Systems*. Princeton University Press: Princeton, 1961.
96. T.L. Casavant and John G. Kuhl. Effects of response and stability on scheduling in distributed computing systems. *IEEE Transaction on Software Engineering*, 14(11):1578–1588, November 1988.

97. E.W. Dijkstra. Self-stabilization in spite of distributed control. In *Texts and Monographs in Computer Science*. pp. 41–46: Springer-Verlag: New York, 1982.
98. E.W. Dijkstra. Self-stabilizing systems in spite of distributed control. *Communications of the ACM*, 17(11):643–644, November 1974.
99. B. Awerbuch and G. Varghese. Distributed program checking: A paradigm for building self-stabilizing distributed protocols. In *Proceedings of the 32nd IEEE Annual Symposium on Foundations of Computer Science*, pp. 258–267, Los Alamitos, CA, 1991.
100. K. Mani Chandy and L. Lamport. Distributed snapshots: Determining global states of distributed systems. *ACM Transactions on Computer Systems*, 3(1):63–75, February 1985.
101. S. Venkatesan and B. Dathan. Testing and debugging distributed programs using global predicates. *IEEE Transaction on Software Engineering*, 21(2):163–177, February 1995.
102. V.K. Garg and B. Waldecker. Detection of strong unstable predicates in distributed programs. Department of Electrical and Computer Engineering, University of Texas, Austin.
103. F.A. Schreiber. Notes on real-time distributed database systems stability. In *Proceedings of the 5th Jerusalem Conference on Information Technology*, pp. 560–564, Los Alamitos, CA, 1990.
104. J.A. Stankovic. Stability and distributed scheduling algorithms. *IEEE Transaction on Software Engineering*, SE-11(10):1141–1152, October 1985.
105. J.F. Meyer. Performability modeling of distributed real-time systems. In G. Iazeolla, P.J. Courtois, and A. Hordijk, ed., *Mathematical Computer Performance and Reliability*, pp. 361–369. Elsevier Science Publishers B.V.: North-Holland, 1984.
106. P. Garg, S. Frolund, A. Shepherd, and N. Venkatasubramanian. Towards distributed applications stability engineering. In *Proceedings of the 5th California Software Symposium*, pp. 528–529, Irvine, CA, March 1995.
107. D. Ferrari. *Computer Systems Performance Evaluation*. Prentice-Hall: Englewood Cliffs, NJ, 1978.
108. Free Software Foundation, Inc. Linux Operating System. Technical report, Cambridge, MA, 1995.