



THE
POWER
TO KNOW.

SAS/OR[®] 12.1 User's Guide Local Search Optimization



The correct bibliographic citation for this manual is as follows: SAS Institute Inc. 2012. *SAS/OR® 12.1 User's Guide: Local Search Optimization*. Cary, NC: SAS Institute Inc.

SAS/OR® 12.1 User's Guide: Local Search Optimization

Copyright © 2012, SAS Institute Inc., Cary, NC, USA

All rights reserved. Produced in the United States of America.

For a hard-copy book: No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, or otherwise, without the prior written permission of the publisher, SAS Institute Inc.

For a Web download or e-book: Your use of this publication shall be governed by the terms established by the vendor at the time you acquire this publication.

The scanning, uploading, and distribution of this book via the Internet or any other means without the permission of the publisher is illegal and punishable by law. Please purchase only authorized electronic editions and do not participate in or encourage electronic piracy of copyrighted materials. Your support of others' rights is appreciated.

U.S. Government Restricted Rights Notice: Use, duplication, or disclosure of this software and related documentation by the U.S. government is subject to the Agreement with SAS Institute and the restrictions set forth in FAR 52.227-19, Commercial Computer Software-Restricted Rights (June 1987).

SAS Institute Inc., SAS Campus Drive, Cary, North Carolina 27513.

Electronic book 1, August 2012

SAS® Publishing provides a complete selection of books and electronic products to help customers use SAS software to its fullest potential. For more information about our e-books, e-learning products, CDs, and hard-copy books, visit the SAS Publishing Web site at support.sas.com/publishing or call 1-800-727-3228.

SAS® and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration.

Other brand and product names are registered trademarks or trademarks of their respective companies.

Contents

Chapter 1.	What's New in SAS/OR 12.1	1
Chapter 2.	Using This Book	9
Chapter 3.	The GA Procedure	15

Subject Index	89
----------------------	-----------

Syntax Index	90
---------------------	-----------

Credits

Documentation

Writing	Joe Hutchinson
Editing	Anne Baxter, Virginia Clark, Ed Huddleston
Documentation Support	Tim Arnold, Remya Chandran, Melanie Gratton, Michelle Opp, Daniel Underwood
Technical Review	Feng Chen, Gehan A. Corea, Tao Huang, Radhika V. Kulkarni

Software

PROC GA	Joe Hutchinson
---------	----------------

Support Groups

Software Testing	Yong Wang, Feng Chen, Tao Huang
Technical Support	Tonya Chapman

Chapter 1

What's New in SAS/OR 12.1

Contents

Overview	1
Highlights of Enhancements in SAS/OR 12.1	1
The CLP Procedure	2
The DTREE, GANTT, and NETDRAW Procedures	3
Supporting Technologies for Optimization	3
PROC OPTMODEL: Nonlinear Optimization	3
Linear Optimization with PROC OPTMODEL and PROC OPTLP	4
Mixed Integer Linear Optimization with PROC OPTMODEL and PROC OPTMILP	4
The Decomposition Algorithm	4
Setting the Cutting Plane Strategy	5
Conflict Search	5
PROC OPTMILP: Option Tuning	6
PROC OPTMODEL: The SUBMIT Block	6
Network Optimization with PROC OPTNET	6
SAS Simulation Studio 12.1	7

Overview

SAS/OR 12.1 delivers a broad range of new capabilities and enhanced features, encompassing optimization, constraint programming, and discrete-event simulation. SAS/OR 12.1 enhancements significantly improve performance and expand your tool set for building, analyzing, and solving operations research models.

In previous years, SAS/OR® software was updated only with new releases of Base SAS® software, but this is no longer the case. This means that SAS/OR software can be released to customers when enhancements are ready, and the goal is to update SAS/OR every 12 to 18 months. To mark this newfound independence, the release numbering scheme for SAS/OR is changing with this release. This new numbering scheme will be maintained when new versions of Base SAS and SAS/OR ship at the same time. For example, when Base SAS 9.4 is released, SAS/OR 13.1 will be released.

Highlights of Enhancements in SAS/OR 12.1

Highlights of the SAS/OR enhancements include the following:

- multithreading is used to improve performance in these three areas:

- PROC OPTMODEL problem generation
- multistart for nonlinear optimization
- option tuning for mixed integer linear optimization
- concurrent solve capability (experimental) for linear programming (LP) and nonlinear programming (NLP)
- improvements to all simplex LP algorithms and mixed integer linear programming (MILP) solver
- new decomposition (DECOMP) algorithm for LP and MILP
- new option for controlling MILP cutting plane strategy
- new conflict search capability for MILP solver
- option tuning for PROC OPTMILP
- new procedure, PROC OPTNET, for network optimization and analysis
- new SUBMIT block for invoking SAS code within PROC OPTMODEL
- SAS Simulation Studio improvements:
 - one-click connection of remote blocks in large models
 - autoscrolling for navigating large models
 - new search capability for block types and label content
 - alternative Experiment window configuration for large experiments
 - selective animation capability
 - new submodel component (experimental)

The CLP Procedure

In SAS/OR 12.1, the CLP procedure adds two classes of constraints that expand its capabilities and can accelerate its solution process. The LEXICO statement imposes a lexicographic ordering between pairs of variable lists. Lexicographic order is essentially analogous to alphabetical order but expands the concept to include numeric values. One vector (list) of values is lexicographically less than another if the corresponding elements are equal up to a certain point and immediately after that point the next element of the first vector is numerically less than the second. Lexicographic ordering can be useful in eliminating certain types of symmetry that can arise among solutions to constraint satisfaction problems (CSPs). Imposing a lexicographic ordering eliminates many of the mutually symmetric solutions, reducing the number of permissible solutions to the problem and in turn shortening the solution process.

Another constraint class that is added to PROC CLP for SAS/OR 12.1 is the bin-packing constraint, imposed via the PACK statement. A bin-packing constraint directs that a specified number of items must be placed into a specified number of bins, subject to the capacities (expressed in numbers of items) of the bins. The PACK statement provides a compact way to express such constraints, which can often be useful components of larger CSPs or optimization problems.

The DTREE, GANTT, and NETDRAW Procedures

In SAS/OR 12.1 the DTREE, GANTT, and NETDRAW procedures each add procedure-specific graph styles that control fonts, line colors, bar and node fill colors, and background images.

Supporting Technologies for Optimization

The underlying improvements in optimization in SAS/OR 12.1 are chiefly related to multithreading, which denotes the use of multiple computational cores to enable computations to be executed in parallel rather than serially. Multithreading can provide dramatic performance improvements for optimization because these underlying computations are performed many times in the course of an optimization process.

The underlying linear algebra operations for the linear, quadratic, and nonlinear interior point optimization algorithms are now multithreaded. The LP, QP, and NLP solvers can be used by PROC OPTMODEL, PROC OPTLP, and PROC OPTQP in SAS/OR. For nonlinear optimization with PROC OPTMODEL, the evaluation of nonlinear functions is multithreaded for improved performance.

Finally, the process of creating an optimization model from PROC OPTMODEL statements has been multithreaded. PROC OPTMODEL contains powerful declarative and programming statements and is adept at enabling data-driven definition of optimization models, with the result that a rather small section of PROC OPTMODEL code can create a very large optimization model when it is executed. Multithreading can dramatically shorten the time that is needed to create an optimization model.

In SAS/OR 12.1 you can use the NTHREADS= option in the PERFORMANCE statement in PROC OPTMODEL and other SAS/OR optimization procedures to specify the number of cores to be used. Otherwise, SAS detects the number of cores available and uses them.

PROC OPTMODEL: Nonlinear Optimization

The nonlinear optimization solver that PROC OPTMODEL uses builds on the introduction of multithreading for its two most significant improvements in SAS/OR 12.1. First, in addition to the nonlinear solver options ALGORITHM=ACTIVESET and ALGORITHM=INTERIORPOINT, SAS/OR 12.1 introduces the ALGORITHM=CONCURRENT option (experimental), with which you can invoke both the active set and interior point algorithms for the specified problem, running in parallel on separate threads. The solution process terminates when either of the algorithms terminates. For repeated solves of a number of similarly structured problems or simply for problems for which the best algorithm isn't readily apparent, ALGORITHM=CONCURRENT should prove useful and illuminating.

Second, multithreading is central to the nonlinear optimization solver's enhanced multistart capability, which now takes advantage of multiple threads to execute optimizations from multiple starting points in parallel. The multistart capability is essential for problems that feature nonconvex nonlinear functions in either or both of the objective and the constraints because such problems might have multiple locally optimal points. Starting optimization from several different starting points helps to overcome this difficulty, and multithreading this process helps to ensure that the overall optimization process runs as fast as possible.

Linear Optimization with PROC OPTMODEL and PROC OPTLP

Extensive improvements to the primal and dual simplex linear optimization algorithms produce better performance and better integration with the crossover algorithm, which converts solutions that are found by the interior point algorithm into more usable basic optimal solutions. The crossover algorithm itself has undergone extensive enhancements that improve its speed and stability.

Paralleling developments in nonlinear optimization, SAS/OR 12.1 linear optimization introduces a concurrent algorithm, invoked with the `ALGORITHM=CONCURRENT` option, in the `SOLVE WITH LP` statement for PROC OPTMODEL or in the PROC OPTLP statement. The concurrent LP algorithm runs a selection of linear optimization algorithms in parallel on different threads, with settings to suit the problem at hand. The optimization process terminates when the first algorithm identifies an optimal solution. As with nonlinear optimization, the concurrent LP algorithm has the potential to produce significant reductions in the time needed to solve challenging problems and to provide insights that are useful when you solve a large number of similarly structured problems.

Mixed Integer Linear Optimization with PROC OPTMODEL and PROC OPTMILP

Mixed integer linear optimization in SAS/OR 12.1 builds on and extends the advances in linear optimization. Overall, solver speed has increased by over 50% (on a library of test problems) compared to SAS/OR 9.3. The branch-and-bound algorithm has approximately doubled its ability to evaluate and solve component linear optimization problems (which are referred to as nodes in the branch-and-bound tree). These improvements have significantly reduced solution time for difficult problems.

The Decomposition Algorithm

The most fundamental change to both linear and mixed integer linear optimization in SAS/OR 12.1 is the addition of the decomposition (DECOMP) algorithm, which is invoked with a specialized set of options in the `SOLVE WITH LP` and `SOLVE WITH MILP` statements for PROC OPTMODEL or in the `DECOMP` statement for PROC OPTLP and PROC OPTMILP. For many linear and mixed integer linear optimization problems, most of the constraints apply only to a small set of decision variables. Typically there are many such sets of constraints, complemented by a small set of linking constraints that apply to all or most of the decision variables. Optimization problems with these characteristics are said to have a “block-angular” structure, because it is easy to arrange the rows of the constraint matrix so that the nonzero values, which correspond to the local sets of constraints, appear as blocks along the main diagonal.

The DECOMP algorithm exploits this structure, decomposing the overall optimization problem into a set of component problems that can be solved in parallel on separate computational threads. The algorithm repeatedly solves these component problems and then cycles back to the overall problem to update key information that is used the next time the component problems are solved. This process repeats until it produces a solution to the complete problem, with the linking constraints present. The combination of parallelized solving of the component problems and the iterative coordination with the solution of the overall

problem can greatly reduce solution time for problems that were formerly regarded as too time-consuming to solve practically.

To use the DECOMP algorithm, you must either manually or automatically identify the blocks of the constraint matrix that correspond to component problems. The METHOD= option controls the means by which blocks are identified. METHOD=USER enables you to specify the blocks yourself, using the .block suffix to declare blocks. This is by far the most common method of defining blocks. If your problem has a significant or dominant network structure, you can use METHOD=NETWORK to identify the blocks in the problem automatically. Finally, if no linking constraints are present in your problem, then METHOD=AUTO identifies the blocks automatically.

The DECOMP algorithm uses a number of detailed options that specify how the solution processes for the component problems and the overall problem are configured and how they coordinate with each other. You can also specify the number of computational threads to make available for processing component problems and the level of detail in the information to appear in the SAS log. Options specific to the linear and mixed integer linear solvers that are used by the DECOMP algorithm are largely identical to those for the respective solvers.

Setting the Cutting Plane Strategy

Cutting planes are a major component of the mixed integer linear optimization solver, accelerating its progress by removing fractional (not integer feasible) solutions. SAS/OR 12.1 adds the CUTSTRATEGY= option in the PROC OPTMILP statement and in the SOLVE WITH MILP statement for PROC OPTMODEL, enabling you to determine the aggressiveness of your overall cutting plane strategy. This option complements the individual cut class controls (CUTCLQUE=, CUTGOMORY=, CUTMIR=, and so on), with which you can enable or disable certain cut types, and the ALLCUTS= option, which enables or disables all cutting planes. In contrast, the CUTSTRATEGY= option controls cuts at a higher level, creating a profile for cutting plane use. As the cut strategy becomes more aggressive, more effort is directed toward creating cutting planes and more cutting planes are applied. The available values of the CUTSTRATEGY= option are AUTOMATIC, BASIC, MODERATE, and AGGRESSIVE; the default is AUTOMATIC. The precise cutting plane strategy that corresponds to each of these settings can vary from problem to problem, because the strategy is also tuned to suit the problem at hand.

Conflict Search

Another means of accelerating the solution process for mixed integer linear optimization takes information from infeasible linear optimization problems that are encountered during an initial exploratory phase of the branch-and-bound process. This information is analyzed and ultimately is used to help the branch-and-bound process avoid combinations of decision variable values that are known to lead to infeasibility. This approach, known as conflict analysis or conflict search, influences presolve operations on branch-and-bound nodes, cutting planes, computation of decision variable bounds, and branching. Although the approach is complex, its application in SAS/OR 12.1 is straightforward. The CONFLICTSEARCH= option in the PROC OPTMILP statement or the SOLVE WITH MILP statement in PROC OPTMODEL enables you to specify the level of conflict search to be performed. The available values for the CONFLICTSEARCH=

option are NONE, AUTOMATIC, MODERATE, and AGGRESSIVE. A more aggressive search strategy explores more branch-and-bound nodes initially before the branch-and-bound algorithm is restarted with information from infeasible nodes included. The default value is AUTOMATIC, which enables the solver to choose the search strategy.

PROC OPTMILP: Option Tuning

The final SAS/OR 12.1 improvement to the mixed integer linear optimization solver is option tuning, which helps you determine the best option settings for PROC OPTMILP. There are many options and settings available, including controls on the presolve process, branching, heuristics, and cutting planes. The TUNER statement enables you to investigate the effects of the many possible combinations of option settings on solver performance and determine which should perform best. The PROBLEMS= option enables you to submit several problems for tuning at once. The OPTIONMODE= option specifies the options to be tuned. OPTIONMODE=USER indicates that you will supply a set of options and initial values via the OPTIONVALUES= data set, OPTIONMODE=AUTO (the default) tunes a small set of predetermined options, and OPTIONMODE=FULL tunes a much more extensive option set.

Option tuning starts by using an initial set of option values to solve the problem. The problem is solved repeatedly with different option values, with a local search algorithm to guide the choices. When the tuning process terminates, the best option values are output to a data set specified by the SUMMARY= option. You can control the amount of time used by this process by specifying the MAXTIME= option. You can multithread this process by using the NTHREADS= option in the PERFORMANCE statement for PROC OPTMILP, permitting analyses of various settings to occur simultaneously.

PROC OPTMODEL: The SUBMIT Block

In SAS/OR 12.1, PROC OPTMODEL adds the ability to execute other SAS code nested inside PROC OPTMODEL syntax. This code is executed immediately after the preceding PROC OPTMODEL syntax and before the syntax that follows. Thus you can use the SUBMIT block to, for example, invoke other SAS procedures to perform analyses, to display results, or for other purposes, as an integral part of the process of creating and solving an optimization model with PROC OPTMODEL. This addition makes it even easier to integrate the operation of PROC OPTMODEL with other SAS capabilities.

To create a SUBMIT block, use a SUBMIT statement (which must appear on a line by itself) followed by the SAS code to be executed, and terminate the SUBMIT block with an ENDSUBMIT statement (which also must appear on a line by itself). The SUBMIT statement enables you to pass PROC OPTMODEL parameters, constants, and evaluated expressions to the SAS code as macro variables.

Network Optimization with PROC OPTNET

PROC OPTNET, new in SAS/OR 12.1, provides several algorithms for investigating the characteristics of networks and solving network-oriented optimization problems. A network, sometimes referred to as a graph,

consists of a set of nodes that are connected by a set of arcs, edges, or links. There are many applications of network structures in real-world problems, including supply chain analysis, communications, transportation, and utilities problems. PROC OPTNET addresses the following classes of network problems:

- biconnected components
- maximal cliques
- connected components
- cycle detection
- weighted matching
- minimum-cost network flow
- minimum cut
- minimum spanning tree
- shortest path
- transitive closure
- traveling salesman

PROC OPTNET syntax provides a dedicated statement for each problem class in the preceding list.

The formats of PROC OPTNET input data sets are designed to fit network-structured data, easing the process of specifying network-oriented problems. The underlying algorithms are highly efficient and can successfully address problems of varying levels of detail and scale. PROC OPTNET is a logical destination for users who are migrating from some of the legacy optimization procedures in SAS/OR. Former users of PROC NETFLOW can turn to PROC OPTNET to solve shortest-path and minimum-cost network flow problems, and former users of PROC ASSIGN can instead use the LINEAR_ASSIGNMENT statement in PROC OPTNET to solve assignment problems.

SAS Simulation Studio 12.1

SAS Simulation Studio 12.1, a component of SAS/OR 12.1 for Windows environments, adds several features that improve your ability to build, explore, and work with large, complex discrete-event simulation models. Large models present a number of challenges to a graphical user interface such as that of SAS Simulation Studio. Connection of model components, navigation within a model, identification of objects or areas of interest, and management of different levels of modeling are all tasks that can become more difficult as the model size grows significantly beyond what can be displayed on one screen. An indirect effect of model growth is an increased number of factors and responses that are needed to parameterize and investigate the performance of the system being modeled.

Improvements in SAS Simulation Studio 12.1 address each of these issues. In SAS Simulation Studio, you connect blocks by dragging the cursor to create links between output and input ports on regular blocks and Connector blocks. SAS Simulation Studio 12.1 automatically scrolls the display of the Model window as

you drag the link that is being created from its origin to its destination, thus enabling you to create a link between two blocks that are located far apart (additionally you can connect any two blocks by clicking on the OutEntity port of the first block and then clicking on the InEntity port of the second block). Automatic scrolling also enables you to navigate a large model more easily. To move to a new area in the Model window, you can simply hold down the left mouse button and drag the visible region of the model to the desired area. This works for simple navigation and for moving a block to a new, remote location in the model.

SAS Simulation Studio 12.1 also enables you to search among the blocks in a model and identify the blocks that have a specified type, a certain character string in their label, or both. From the listing of identified blocks, you can open the Properties dialog box for each identified block and edit its settings. Thus, if you can identify a set of blocks that need similar updates, then you can make these updates without manually searching through the model for qualifying blocks and editing them individually. For very large models, this capability not only makes the update process easier but also makes it more thorough because you can identify qualifying blocks centrally.

When you design experiments for large simulation models, you often need a large number of factors to parameterize the model and a large number of responses to track system performance in sufficient detail. This was a challenge prior to SAS Simulation Studio 12.1 because the Experiment window displayed factors and responses in the header row of a table, with design points and their replications' results displayed in the rows below. A very large number of factors and responses did not fit on one screen in this display scheme, and you had to scroll across the Experiment window to view all of them.

SAS Simulation Studio 12.1 provides you with two alternative configurations for the Experiment window. The Design Matrix tab presents the tabular layout described earlier. The Design Point tab presents each design point in its own display. Factors and responses (summarized over replications) are displayed in separate tables, each with the factor or response names appearing in one column and the respective values in a second column. This layout enables a large number of factors and responses to be displayed. Response values for each replication of the design point can be displayed in a separate window.

SAS Simulation Studio 12.1 enhances its multilevel model management features by introducing the submodel component (experimental). Like the compound block, the submodel encapsulates a group of SAS Simulation Studio blocks and their connections, but the submodel outpaces the compound block in some important ways. The submodel, when expanded, opens in its own window. This means a submodel in its collapsed form can be placed close to other blocks in the Model window without requiring space for its expanded form (as is needed for compound blocks). The most important property of the submodel is its ability to be copied and instantiated in several locations simultaneously, whether in the same model, in different models in the same project, or in different projects. Each such instance is a direct reference to the original submodel, not a disconnected copy. Thus you can edit the submodel by editing any of its instances; changes that are made to any instance are propagated to all current and future instances of the submodel. This feature enables you to maintain consistency across your models and projects.

Finally, SAS Simulation Studio 12.1 introduces powerful new animation controls that should prove highly useful in debugging simulation models. In the past, animation could be switched on or off and its speed controlled, but these choices were made for the entire model. If you needed to animate a particular segment of the model, perhaps during a specific time span for the simulation clock, you had to focus your attention on that area and pay special attention when the time period of interest arrived. In SAS Simulation Studio 12.1 you can select both the area of the model to animate (by selecting a block or a compound block) and the time period over which animation should occur (by specifying the start and end times for animation). You can also control simulation speed for each such selection. Multiple selections are supported so that you can choose to animate several areas of the model, each during its defined time period and at its chosen speed.

Chapter 2

Using This Book

Contents	
Purpose	9
Organization	9
Typographical Conventions	10
Conventions for Examples	11
Accessing the SAS/OR Sample Library	11
Online Documentation	11
Additional Documentation for SAS/OR Software	12

Purpose

SAS/OR User’s Guide: Local Search Optimization provides a complete reference for the local search procedure in SAS/OR software. This book serves as the primary documentation for the GA procedure, which uses genetic algorithms to solve optimization problems

“Using This Book” describes the organization of this book and the conventions used in the text and example code. To gain full benefit from using this book, you should familiarize yourself with the information presented in this section and refer to it when needed. The section “[Additional Documentation for SAS/OR Software](#)” on page 12 refers to other documents that contain related information.

Organization

[Chapter 3](#) describes the GA procedure. The procedure description is self-contained; you need to be familiar with only the basic features of the SAS System and SAS terminology to use it. The following list summarizes the types of information provided for the GA procedure:

- Overview** provides a general description of what the procedure does. It outlines major capabilities of the procedure and lists all input and output data sets that are used with it.
- Getting Started** illustrates simple uses of the procedure using a few short examples. It provides introductory *hands-on* information for the procedure.

- Syntax** constitutes the major reference section for the syntax of the procedure. First, the statement syntax is summarized. Next, a functional summary table lists all the statements and options in the procedure, classified by function. In addition, the online version includes a Dictionary of Options, which provides an alphabetical list of all options in the GA procedure. The PROC GA statement is then described.
- Details** describes the features of the procedure, including algorithmic details and computational methods. It also explains how the various options interact with each other. This section describes input and output data sets in greater detail, with definitions of the output variables, and explains the format of printed output, if any.
- Examples** consists of examples that are designed to illustrate the use of the procedure. Each example includes a description of the problem and lists the options that are highlighted by the example. The example shows the data and the SAS statements needed, and includes the output produced. You can duplicate the examples by copying the statements and data and running the SAS program. The SAS Sample Library contains the code used to run the examples shown in this book; consult your SAS Software representative for specific information about the Sample Library.
- References** lists references that are relevant to the chapter.

Typographical Conventions

The printed version of *SAS/OR User's Guide: Local Search Optimization* uses various type styles, as explained by the following list:

- | | |
|-----------------|--|
| roman | is the standard type style used for most text. |
| UPPERCASE ROMAN | is used for SAS statements, options, and other SAS language elements when they appear in the text. However, you can enter these elements in your own SAS code in lowercase, uppercase, or a mixture of the two. This style is also used for identifying arguments and values (in the syntax specifications) that are literals (for example, to denote valid keywords for a specific option). |

UPPERCASE BOLD	is used in the “Syntax” section to identify SAS keywords such as the names of procedures, statements, and options.
VariableName	is used for the names of SAS variables and data sets when they appear in the text.
<i>oblique</i>	is used to indicate an option variable for which you must supply a value (for example, DUPLICATE= <i>dup</i> indicates that you must supply a value for <i>dup</i>).
<i>italic</i>	is used for terms that are defined in the text, for emphasis, and for publication titles.
monospace	is used to show examples of SAS statements. In most cases, this book uses lowercase type for SAS code. You can enter your own SAS code in lowercase, uppercase, or a mixture of the two.

Conventions for Examples

Most of the output shown in this book is produced with the following SAS System options:

```
options linesize=80 pagesize=60 nonumber nodate;
```

Accessing the SAS/OR Sample Library

The SAS/OR sample library includes many examples that illustrate the use of SAS/OR software, including the examples used in this documentation. To access these sample programs from the SAS windowing environment, select **Help** from the main menu and then select **Getting Started with SAS Software**. On the **Contents** tab, expand the **Learning to Use SAS, Sample SAS Programs**, and **SAS/OR** items. Then click **Samples**.

Online Documentation

This documentation is available online with the SAS System. To access SAS/OR documentation from the SAS windowing environment, select **Help** from the main menu and then select **SAS Help and Documentation**. On the **Contents** tab, expand the **SAS Products** and **SAS/OR** items. Then expand the book you want to view. You can search the documentation by using the **Search** tab.

You can also access the documentation by going to <http://support.sas.com/documentation>.

Additional Documentation for SAS/OR Software

In addition to *SAS/OR User's Guide: Local Search Optimization*, you might find the following documents helpful when using SAS/OR software:

SAS/OR User's Guide: Bill of Material Processing

provides documentation for the BOM procedure and all bill of material postprocessing SAS macros. The BOM procedure and SAS macros provide the ability to generate different reports and to perform several transactions to maintain and update bills of material.

SAS/OR User's Guide: Constraint Programming

provides documentation for the constraint programming procedure in SAS/OR software. This book serves as the primary documentation for the CLP procedure.

SAS/OR User's Guide: Mathematical Programming

provides documentation for the mathematical programming procedures in SAS/OR software. This book serves as the primary documentation for the OPTLP, OPTMILP, OPTMODEL, and OPTQP procedures, the various solvers called by the OPTMODEL procedure, and the MPS-format SAS data set specification.

SAS/OR User's Guide: Mathematical Programming Examples

supplements the *SAS/OR User's Guide: Mathematical Programming* with additional examples that demonstrate best practices for building and solving linear programming, mixed integer linear programming, and quadratic programming problems. The problem statements are reproduced with permission from the book *Model Building in Mathematical Programming* by H. Paul Williams.

SAS/OR User's Guide: Mathematical Programming Legacy Procedures

provides documentation for the older mathematical programming procedures in SAS/OR software. This book serves as the primary documentation for the INTPOINT, LP, NETFLOW, and NLP procedures. Guidelines are also provided on migrating from these older procedures to the newer OPTMODEL family of procedures.

SAS/OR User's Guide: Network Optimization Algorithms

provides documentation for a set of algorithms that can be used to investigate the characteristics of networks and to solve network-oriented optimization problems. This book also documents PROC OPTNET, which invokes these algorithms and provides network-structured formats for input and output data.

SAS/OR User's Guide: Project Management

provides documentation for the project management procedures in SAS/OR software. This book serves as the primary documentation for the CPM, DTREE, GANTT, NETDRAW, and PM procedures, the earned value management macros, the Microsoft Project conversion macros, and the PROJMAN application.

SAS/OR Software: Project Management Examples, Version 6

contains a series of examples that illustrate how to use SAS/OR software to manage projects. Each chapter contains a complete project management scenario and describes how to use PROC GANTT, PROC CPM, and PROC NETDRAW, in addition to other reporting and graphing procedures in the SAS System, to perform the necessary project management tasks.

SAS Simulation Studio: User's Guide

provides documentation for using SAS Simulation Studio, a graphical application for creating and working with discrete-event simulation models. This book describes in detail how to build and run simulation models and how to interact with SAS software for analysis and with JMP software for experimental design and analysis.

Chapter 3

The GA Procedure

Contents

Overview: GA Procedure	16
Getting Started: GA Procedure	18
Initializing the Problem Data	19
Choosing the Problem Encoding	21
Setting the Objective Function	22
Controlling the Selection Process	22
Setting Crossover Parameters	23
Setting Mutation Parameters	24
Creating the Initial Generation	24
Monitoring Progress and Reporting Results	25
A Simple Example	25
Syntax: GA Procedure	27
PROC GA Statement	28
ContinueFor Call	30
Cross Call	31
Dynamic_array Call	31
EvaluateLC Call	32
GetDimensions Call	33
GetObjValues Call	33
GetSolutions Call	33
Initialize Call	34
MarkPareto Call	35
Mutate Call	36
Objective Function	37
PackBits Call	38
Programming Statements	38
ReadChild Call	40
ReadCompare Call	41
ReadMember Call	41
ReadParent Call	42
ReEvaluate Call	42
SetBounds Call	42
SetCompareRoutine Call	43
SetCross Call	43
SetCrossProb Call	44
SetCrossRoutine Call	45

SetElite Call	45
SetEncoding Call	45
SetFinalize Call	46
SetMut Call	46
SetMutProb Call	47
SetMutRoutine Call	47
SetObj Call	48
SetObjFunc Call	48
SetProperty Call	49
SetSel Call	49
SetUpdateRoutine Call	50
ShellSort Call	50
Shuffle Call	51
UnpackBits Function	51
UpdateSolutions Call	52
WriteChild Call	52
WriteMember Call	52
Details: GA Procedure	53
Using Multisegment Encoding	53
Using Standard Genetic Operators and Objective Functions	54
Defining a User Fitness Comparison Routine	63
Defining User Genetic Operators	64
Defining a User Update Routine	67
Defining an Objective Function	68
Defining a User Initialization Routine	69
Specifying the Selection Strategy	71
Incorporating Heuristics and Local Optimizations	72
Handling Constraints	72
Optimizing Multiple Objectives	73
Examples: GA Procedure	74
Example 3.1: Traveling Salesman Problem with Local Optimization	74
Example 3.2: Nonlinear Objective with Constraints Using Repair Mechanism	78
Example 3.3: Quadratic Objective with Linear Constraints	81
References	88

Overview: GA Procedure

Genetic algorithms are a family of local search algorithms that seek optimal solutions to problems by applying the principles of natural selection and evolution. Genetic algorithms can be applied to almost any optimization problem and are especially useful for problems where other calculus-based techniques do not work, such as when the objective function has many local optima, when it is not differentiable or continuous, or when solution elements are constrained to be integers or sequences. In most cases genetic algorithms

require more computation than specialized techniques that take advantage of specific problem structures or characteristics. However, for optimization problems with no such techniques available, genetic algorithms provide a robust general method of solution.

In general, genetic algorithms use some variation of the following procedure to search for an optimal solution:

<i>initialization:</i>	An initial population of solutions is randomly generated, and the objective function is evaluated for each member of this initial generation.
<i>selection:</i>	Individual members of the current generation are chosen stochastically either to parent the next generation or to be passed on to it, such that those members who are the fittest are more likely to be selected. A solution's fitness is based on its objective value, with better objective values reflecting higher fitness.
<i>crossover:</i>	Some of the selected solutions are passed to a crossover operator. The crossover operator combines two or more parents to produce new offspring solutions for the next generation. The crossover operator tends to produce new offspring that retain the common characteristics of the parent solutions, while combining the other traits in new ways. In this way new areas of the search space are explored, hopefully while retaining optimal solution characteristics.
<i>mutation:</i>	Some of the next-generation solutions are passed to a mutation operator, which introduces random variations in the solutions. The purpose of the mutation operator is to ensure that the solution space is adequately searched to prevent premature convergence to a local optimum.
<i>repeat:</i>	The current generation of solutions is replaced by the new generation. If the stopping criterion is not satisfied, the process returns to the <i>selection</i> phase.

The crossover and mutation operators are commonly called *genetic operators*. Selection and crossover distinguish genetic algorithms from a purely random search and direct the algorithm toward finding an optimum. Mutation is designed to ensure diversity in the search to prevent premature convergence to a local optimum.

There are many ways to implement the general strategy just outlined, and it is also possible to combine the genetic algorithm approach with other heuristic solution improvement techniques. In the traditional genetic algorithm, the solutions space is composed of bit strings, mapped to an objective function, and the genetic operators are modeled after biological processes. Although there is a theoretical foundation for the convergence of genetic algorithms formulated in this way, in practice most problems do not fit naturally into this paradigm. Modern research has shown that optimizations can be set up by using the natural solution domain (for example, a real vector or integer sequence) and applying crossover and mutation operators analogous to the traditional genetic operators, but more appropriate to the natural formulation of the problem. This is the approach, sometimes called *evolutionary computing*, taken in the GA procedure. It enables you to model your problem by using a variety of solution forms including sequences, integer or real

vectors, Boolean encodings, and combinations of these. The GA procedure also provides you with a choice of genetic operators appropriate for these encodings, while permitting you to write your own.

The GA procedure enables you to implement the basic genetic algorithm by default, and to employ other advanced techniques to handle constraints, accelerate convergence, and perform multiobjective optimizations. These advanced techniques are discussed in the section “[Details: GA Procedure](#)” on page 53.

Although genetic algorithms have been demonstrated to work well for a variety of problems, there is no guarantee of convergence to a global optimum. Also, the convergence of genetic algorithms can be sensitive to the choice of genetic operators, mutation probability, and selection criteria, so that some initial experimentation and fine-tuning of these parameters is often required.

Getting Started: GA Procedure

The optimization problem is described by using programming statements, which initialize problem data and specify the objective, genetic operators, and other optimization parameters. The programming statements are executed once, and are followed by a RUN statement to begin the optimization process. The GA procedure enables you to define subroutines and designate them to be called during the optimization process to calculate objective functions, perform genetic mutation or crossover operations, or monitor and control the optimization. All variables created within a subroutine are local to that routine; to access a global variable defined within the GA procedure, the subroutine must have a parameter with the same name as the variable.

To set up a genetic algorithm optimization, your program needs to perform the following steps:

1. Initialize the problem data, such as cost coefficients and parameter limits.
2. Specify five basic optimization parameters:
 - *Encoding*: the general structure and form of the solution
 - *Objective*: the function to be optimized
 - *Selection*: how members of the current solution generation are chosen to propagate the next generation
 - *Crossover*: how the attributes of parent solutions are combined to produce new offspring solutions
 - *Mutation*: how random variation is introduced into the new offspring solutions to maintain genetic diversity
3. Generate a population of solutions for the initial generation.
4. Control the execution of the algorithm and record your results.

The following sections discuss each of these items in detail.

Initializing the Problem Data

The GA procedure offers great flexibility in how you initialize the problem data. Either you can read data from SAS data sets that are created from other SAS procedures and DATA steps, or you can initialize the data with programming statements.

In the PROC GA statement, you can specify up to five data sets to be read with the DATA n = option, where n is a number from 1 to 5, that can be used to initialize parameters and data vectors applicable to the optimization problem. For example, weights and rewards for a knapsack problem could be stored in the variables WEIGHT and REWARD in a SAS data set. If you specify the data set with a DATA1= option, the arrays WEIGHT and REWARD are initialized at the start of the procedure and are available for computing the objective function and evaluating the constraints with program statements. You could store the number of items and weight limit constraint in another data set, as illustrated in the sample programming statements that follow:

```
data input1;
    input weight reward;
    datalines;
1      5
2      3
4      7
1      2
8      3
6      9
2      6
4      3
...
;

data input2;
    input nitems limit;
    datalines;
10  20
;

proc ga data1 = input1 /* creates arrays weight and reward */
      data2 = input2; /* creates variables nitems and limit */

function objective( selected[*], reward[*], nitems);
    array x[1] /nosym;
    call dynamic_array(x, nitems);
    call ReadMember(selected,1,x);
    obj = 0;
    do i=1 to nitems;
        obj = obj + reward[x[i]];
    end;
    return(obj);
endsub;

[Other statements follow]
```

With these statements, the DATA1= option first establishes the arrays `weight` and `reward` from the data set `input1`, and the DATA2= option causes the variables `nitems` and `limit` to be created and initialized from the data set `input2`. The `reward` array and the `nitems` variable are then used in the objective function.

For convenience in initializing two-dimensional data such as matrices, the GA procedure provides you with the `MATRIXn=` option, where n is a number from 1 to 5. A two-dimensional array is created within the GA procedure with the same name as the option, containing the numeric data in the specified data set. For example, a table of distances between cities for a traveling salesman problem could be stored as a SAS data set, and a `MATRIX1=` option specifying that data set would cause a two-dimensional array named `MATRIX1` to be created containing the data at the start of the GA procedure. This is illustrated in the following program:

```
data distance;
    input d1-d10;
    datalines;
0 5 3 1 2 ...
5 0 4 2 6 ...
3 4 0 1 3 ...
...
;

proc ga matrix1 = distance;
ncities = 10;
call SetEncoding('S10');
call SetObj('TSP', 'distances', matrix1);

[Other statements follow]
```

In this example, the data set `distance` is used to create a two-dimensional array `matrix1`, where `matrix1[i, j]` is the distance from city i to city j . The GA procedure provides a simple traveling salesman Problem (TSP) objective function, which is specified by the user with the `SetObj` call. The distances between locations are specified with the `distances` property of the TSP objective, which is set in the call to be `matrix1`. Note that when a `MATRIXn=` option is used, the names of variables in the data set are not transferred to the GA procedure as they are with a `DATA` option; only the numeric data are transferred.

You can also initialize problem data with programming statements. The programming statements in the GA procedure are executed before the optimization process begins. The variables created and initialized can be used and modified as the optimization progresses. The programming statement syntax is much like the SAS DATA step, with a few differences (see the section “[Syntax: GA Procedure](#)” on page 27). Special calls are described in the next sections that enable you to specify the objective function and genetic operators, and to monitor and control the optimization process. In the following program, a two-dimensional matrix is set up with programming statements to provide the distances for a 10-city symmetric traveling salesman problem, between locations specified in a SAS data set:

```
data positions;
    input x y;
    datalines;
100 230
50 20
150 100
...
;
```

```

proc ga data1 = positions;

call SetEncoding('S10');
ncities = 10;

array distance[10,10] /nosym;

do i = 1 to ncities;
  do j = 1 to i;
    distance[i,j] = sqrt((x[i]-x[j])**2 + (y[i] - y[j])**2);
    distance[j,i] = distance[i,j];
  end;
end;

call SetObj('TSP', 'distances', distance);

```

In this example, the DATA1= option creates arrays *x* and *y* containing the coordinates of the cities in an *x*-*y* grid, read in from the positions data set. An ARRAY programming statement creates a matrix of distances between cities, and the loops calculate Euclidean distances from the position data. The ARRAY statement is used to create internal data vectors and matrices. It is similar to the ARRAY statement used in the SAS DATA step, but the /NOSYM option is used in this example to set up the array without links to other variables. This option enables the array elements to be indexed more efficiently and the array to be passed efficiently to subroutines. You should use the /NOSYM option whenever you are creating an array that might be passed as an argument to a function or call routine.

Choosing the Problem Encoding

Problem encoding refers to the structure or type of solution space that is to be optimized, such as real-valued fixed-length vectors or integer sequences. The GA procedure offers encoding options appropriate to several types of optimization problems. You specify the problem encoding with a [SetEncoding CALL](#) statement,

```
call SetEncoding('encoding');
```

where the *encoding* string is a letter followed by a number, which specifies the type of encoding and the number of elements. The permitted letters and corresponding types of encoding are as follows:

- R* or *r*: real-valued vector. This type of encoding is used for general non-linear optimization problems.
- I* or *i*: integer-valued vector. This encoding is used for integer-valued problems. The integer size is 32 bits, which imposes a maximum value of 2,147,483,647 and a minimum value of -2,147,483,648 for any element of the vector. The integer vector encoding can also be used to represent bit vectors, where the 0–1 value of each bit in the integer vector is treated as a separate element. An example might be an assignment problem, where the positions within the vector represent different tasks, and the integer values represent different machines or other resources that might be applied to each task.
- B* or *b*: Boolean vector. Each element represents one true/false value.
- S* or *s*: sequence or permutation. In this encoding, each solution is composed of a sequence of integers ranging from 1 to the number of elements, with different solutions distinguished by different

orderings of the elements. This encoding is commonly used for routing problems such as the traveling salesman problem or for scheduling problems.

For example, the following statement specifies a 10-element integer vector encoding:

```
call SetEncoding('I10');
```

For problems where the solution form requires more than one type of encoding, you can specify multiple encodings in the *encoding* string. For example, if you want to optimize the scheduling of 10 tasks and the assignment of resources to each task, you could use a *segmented* encoding, as follows:

```
call SetEncoding('I10S10');
```

Here the I10 (10-element integer vector) is assigned to the first segment, and represents the resource assignment. The S10 (10-element sequence) is assigned to a second segment, and represents the sequence of tasks. The use of segmented encodings is described in the section “[Using Multisegment Encoding](#)” on page 53.

Setting the Objective Function

Before executing a genetic algorithm, you must specify the objective function to be optimized. Either you can define a function in your GA procedure input and designate it to be your objective function with the [SetObjFunc](#) call, or you can specify an objective function that the GA procedure provides with a [SetObj](#) call. The GA procedure currently supports the traveling salesman problem objective.

Controlling the Selection Process

There are two competing factors that need to be balanced in the selection process: *selective pressure* and *genetic diversity*. Selective pressure, the tendency to select only the best members of the current generation to propagate to the next, is required to direct the genetic algorithm to an optimum. Genetic diversity, the maintenance of a diverse solution population, is also required to ensure that the solution space is adequately searched, especially in the earlier stages of the optimization process. Too much selective pressure can lower the genetic diversity so that the global optimum is overlooked and the genetic algorithm converges prematurely. Yet, with too little selective pressure, the genetic algorithm might not converge to an optimum in a reasonable time. A proper balance between the selective pressure and genetic diversity must be maintained for the genetic algorithm to converge in a reasonable time to a global optimum.

The GA procedure uses a standard technique for the selection process commonly known as *tournament selection*. In general, the tournament selection process randomly chooses a group of members from the current population, compares their fitness, and selects the fittest from the group to propagate to the next generation. Tournament selection is one of the fastest selection methods, and it offers good control over the selection pressure.

You can control the selective pressure by specifying the tournament size, the number of members chosen to compete in each tournament. This number should be 2 or greater, with 2 implying the weakest selection

pressure. Tournament sizes from 2 to 10 have been successfully applied to various genetic algorithm optimizations, with sizes over 4 or 5 considered to represent strong selective pressure. This selection option is chosen with the following [SetSel](#) call:

```
call SetSel('tournament', 'size', size);
```

where *size* is the desired tournament size.

For tournament size of 2, you can further weaken the selective pressure by specifying a probability for selecting the most fit solution from the 2 competing solutions. By default this probability is 1, but the GA procedure permits you to set it to a value between 0.5 (equivalent to pure random selection) and 1. This selection option is chosen with the following [SetSel](#) call:

```
call SetSel('duel', 'pbest', bestprob);
```

where *bestprob* is the probability for choosing the most fit solution. This option can prove useful when conventional tournament selection tends to result in premature convergence.

One potential problem with tournament selection is that it does not guarantee that the best solution in the current generation is passed on to the next. To resolve this problem, the GA procedure enables you to specify an *elite* parameter, which ensures that the very best solutions are passed on to the next generation unchanged by mutation or crossover. Use the [SetElite](#) call:

```
call SetElite(elite);
```

where *elite* is an integer greater than or equal to 0. The GA procedure preserves the *elite* best solutions in the current generation and ensures that they are passed on to the next generation unchanged. When writing out the final solution population, the first *elite* members are the best of the generation and are sorted by their fitness, such that the fittest is first. By default, if you do not call [SetElite](#) in your program, an *elite* value of 1 is used. Setting the *elite* parameter to a higher number accelerates the convergence of the genetic algorithm; however, it can also lead to premature convergence before reaching a global optimum, so it should be used with care.

If you do not call [SetSel](#) in your input, then the default behavior for the GA procedure is to use tournament selection with size 2.

Setting Crossover Parameters

There are two crossover parameters that need to be specified: the crossover probability and the crossover operator. Members of the current generation that have passed the selection process either go to the crossover operator or are passed unchanged into the next generation, according to the crossover probability. To set the probability, you use a [SetCrossProb](#) call:

```
call SetCrossProb(prob);
```

where *prob* is a real number between 0 and 1. A value of 1 implies that the crossover operator is always applied, while 0 effectively turns off crossover. If you do not explicitly set the crossover probability with this call, a default value of 0 is used.

The GA procedure enables you to choose your crossover operator from several standard crossover operators appropriate for each type of encoding, or to code your own crossover operator as a subroutine. To specify one of the crossover operators provided by the GA procedure, use a [SetCross](#) call. See the section “[Crossover](#)”

[Operators](#)” on page 55 for more detail on the available operators. To supply your own operator, use a [SetCrossRoutine](#) call:

```
call SetCrossRoutine( 'routine' );
```

where *routine* is the name of your crossover subroutine. See the section “[Defining User Genetic Operators](#)” on page 64 for a description of defining genetic operators with user subroutines. If the crossover probability is greater than 0, you must use a [SetCross](#) or [SetCrossRoutine](#) call to set the crossover operator.

After initialization, you can reset any of the crossover parameters or their properties during the optimization process by calling one of the preceding routines from a user [update routine](#). This makes it possible to adapt the crossover operators as desired in the optimization process.

Setting Mutation Parameters

There are two mutation parameters: the mutation probability and the mutation operator. Members of the next generation are chosen to undergo mutation with the mutation probability you specify. To set the mutation probability, you use a [SetMutProb](#) call:

```
call SetMutProb( prob );
```

where *prob* is a real number between 0 and 1. This probability is usually fairly low (0.05 or less), since mutation tends to slow the convergence of the genetic algorithm. If you do not explicitly set the mutation probability with this call, a default value of 0 is used.

The GA procedure enables you to choose your mutation operator from several standard mutation operators appropriate for each type of encoding, or to code your own mutation operator as a subroutine. To specify one of the mutation operators provided by the GA procedure, use a [SetMut](#) call. See the section “[Mutation Operators](#)” on page 61 for more detail on the available operators. To supply your own operator, use a [SetMutRoutine](#) call:

```
call SetMutRoutine( 'routine' );
```

where *routine* is the name of your mutation subroutine. See the section “[Defining User Genetic Operators](#)” on page 64 for a description of defining genetic operators with user subroutines. If the mutation probability is greater than 0, you must use a [SetMut](#) or [SetMutRoutine](#) call to set the mutation operator.

After initialization, you can reset any of the mutation parameters or their properties during the optimization process by calling one of the preceding routines from a user [update routine](#). This makes it possible to adapt the mutation operators to create more or less diversity as needed in the optimization process.

Creating the Initial Generation

The last step in the initialization for the genetic algorithm optimization is to create the initial solution population, the first generation. The GA procedure provides you with the [Initialize](#) call to accomplish this task. The procedure provides several options for initializing the first population or reinitializing the solution population during the optimization process. For example, you can specify a data set in the `FIRSTGEN=` option of the `PROC GA` statement to be read to populate the initial generation, and use the initialize call:

```
call Initialize( '_dataset_', PopulationSize );
```

Other possible initialization options include generating solutions uniformly distributed over the solution domain, executing a user-defined initialization routine, carrying over a portion of the previous population (for reinitialization), or any combination of those actions. See the section “[Initialize Call](#)” on page 34 for a full explanation of initialization actions.

Monitoring Progress and Reporting Results

The GA procedure enables your program to monitor and alter parameters during the optimization process and record final results.

If a data set is specified in the `LASTGEN=` option of the `PROC GA` statement, then the last generation of solutions is written to the data set. See the section “[PROC GA Statement](#)” on page 28 for a description of the data set created by the `LASTGEN=` option.

You can define a subroutine and designate it to be called at each iteration in an update phase, which occurs after the evaluation phase and before selection, crossover, and mutation. Your subroutine can check solution values and update and store variables you have defined, adjust any of the optimization parameters such as the mutation probability or *elite* value, or check termination criteria and end the optimization process. An update routine can be especially helpful in implementing advanced techniques such as multiobjective optimization. You can specify an update subroutine with a [SetUpdateRoutine call](#):

```
call SetUpdateRoutine('routine');
```

where *routine* is the name of your subroutine to be executed at each iteration.

You can set the maximum number of iterations permitted for the optimization process with the `MAXITER=` option in the `PROC GA` statement. If none is specified, a default value of 500 iterations is used. You can also control the number of iterations dynamically in your program by using the [ContinueFor call](#):

```
call ContinueFor(n);
```

where *n* is the number of iterations to permit beyond the current iteration. A value of 0 ends the optimization process at the current iteration. One common way this call might be used is to include it in the logic of an update subroutine declared in the [SetUpdateRoutine call](#). The update subroutine could check the objective values and end the optimization process when the optimal value of the objective function has not improved for a specific number of iterations. A [ContinueFor call](#) overrides an iteration limit set with the `MAXITER=` option.

To perform post processing of the optimization data, you can use a [SetFinalize call](#) to instruct the GA procedure to call a subroutine you have defined, after the last iteration:

```
call SetFinalize('routine');
```

where *routine* is the name of a subroutine you have defined. Your finalize subroutine could perform some post processing tasks, such as applying heuristics or a local optimization technique to try to improve the final solution.

A Simple Example

The example that follows illustrates the application of genetic algorithms to function optimization over a real-valued domain. It finds the minimum of the Shubert function:

$$\left[\sum_{i=1}^5 i \cos[(i+1)x_1 + i] \right] \left[\sum_{i=1}^5 i \cos[(i+1)x_2 + i] \right]$$

where $-10 \leq x_i \leq 10$ for $i = 1, 2$.

```

proc ga seed = 12 maxiter = 30;
/* the objective function to be optimized */
function shubert(selected[*]);
  array x[2] /nosym;
  call ReadMember(selected,1,x);
  x1 = x[1];
  x2 = x[2];
  sum1 = 0;
  do i = 1 to 5;
    sum1 = sum1 + i * cos((i+1)* x1 + i);
  end;
  sum2 = 0;
  do i = 1 to 5;
    sum2 = sum2 + i * cos((i+1) * x2 + i);
  end;
  result = sum1 * sum2;
  return(result);
endsub;
/* Set the problem encoding */
call SetEncoding('R2');
/* Set upper and lower bounds on the solution components */
array LowerBound[2] /nosym (-10 -10);
array UpperBound[2] /nosym (10 10);
call SetBounds(LowerBound, UpperBound);
/* Set the objective function */
call SetObjFunc('shubert',0);
/* Set the crossover parameters */
call SetCrossProb(0.65);
call SetCross('Heuristic');
/* Set the mutation parameters */
call SetMutProb(0.15);
array del[2] /nosym (0.2 0.2);
call SetMut('Delta','nchange', 1, 'delta',del);
/* Set the selection criteria */
call SetSel('tournament','size', 2);
call SetElite(2);
/* Initialize the first generation, with 120 random solutions */
call Initialize('DEFAULT',120);
/* Now execute the Genetic Algorithm */
run;
quit;

```

At the beginning of the program, the PROC GA statement sets the initial random number seed and sets the maximum number of iterations to 30.

A routine to compute the objective function (`function shubert`) is then defined. This function is called by the GA procedure once for each member of the solution population at each iteration. Note that the GA procedure passes the array `selected` as the first parameter of the function, and the function uses that array

to obtain the selected solution elements with a [ReadMember call](#), which places the solution in the array `x`. The second parameter of the `ReadMember` call is 1, specifying that segment 1 of the solution be returned, which in this case is the only segment. The programming statements that follow compute the value of the objective function and return it to the GA procedure.

After the function definition, the 'R2' passed to the [SetEncoding call](#) specifies that solutions are single-segment, with that segment containing two elements that are real-valued. Next, a lower bound of -10 and an upper bound of 10 are set for each solution element with the [SetBounds call](#). The [SetObjFunc call](#) specifies the previously defined Shubert function as the objective function for the optimization; the second parameter value of 0 indicates that a minimum is desired. The [SetCrossProb call](#) sets the crossover probability to 0.65, which means that, on average, 65% of the solutions passing the selection phase will undergo the crossover operation. The crossover operator is set to the heuristic operator by the [SetCross call](#). Similarly, the mutation probability is set to 0.15 with the [SetMutProb call](#), and the delta operator is set as the mutation operator with the [SetMut call](#). The selection criteria are then set: a conventional tournament of size 2 is specified with [SetSel call](#), and an *elite* value of 2 is specified with the [SetElite call](#). The *elite* value implies that the best two solutions of each generation are always carried over to the next generation unchanged by mutation or crossover. The last step before beginning the optimization is the [Initialize call](#). This call sets the population size at 120, and specifies the default initialization strategy for the first population. For real encoding, this means that an initial population randomly distributed between the upper and lower bounds specified in the [SetBounds call](#) is generated. Finally, when the `RUN` statement is encountered, the GA procedure begins the optimization process. It iterates through 30 generations, as set by the `MAXITER=` option.

The Shubert function has 760 local minima, 18 of which are global minima, with a minimum of -186.73. If you experiment with different random seeds with the `SEED=` option, `PROC GA` generally converges to a different global minimum each time. [Figure 3.1](#) shows the output for the chosen seed.

Figure 3.1 Shubert Function Example Output

PROC GA Optimum Values		
Objective		
-186.7309031		
Solution		
Element	Value	
1	-7.708309818	
2	-0.800371886	

Syntax: GA Procedure

To initialize your data and describe your model, you use programming statements with a syntax similar to the SAS DATA step, augmented with some special function calls to communicate with the genetic algorithm optimizer. Most of the programming statements used in the SAS DATA step can be used in the GA

procedure, and these are described fully in the *SAS Statements: Reference* and Base SAS documentation. Following is an alphabetical list of the statements and special function calls used.

```

PROC GA options ;
  ContinueFor Call ;
  Cross Call ;
  Dynamic_array Call ;
  EvaluateLC Call ;
  GetDimensions Call ;
  GetObjValues Call ;
  GetSolutions Call ;
  Initialize Call ;
  MarkPareto Call ;
  Mutate Call ;
  Objective Call ;
  PackBits Call ;
  Program Statements ;
  ReadChild Call ;
  ReadCompare Call ;
  ReadMember Call ;
  ReadParent Call ;
  ReEvaluate Call ;
  SetBounds Call ;
  SetCross Call ;
  SetCrossProb Call ;
  SetCrossRoutine Call ;
  SetElite Call ;
  SetEncoding Call ;
  SetFinalize Call ;
  SetMut Call ;
  SetMutProb Call ;
  SetMutRoutine Call ;
  SetObj Call ;
  SetObjFunc Call ;
  SetProperty Call ;
  SetSel Call ;
  SetUpdateRoutine Call ;
  ShellSort Call ;
  Shuffle Call ;
  UnpackBits Function ;
  UpdateSolutions Call ;
  WriteChild Call ;
  WriteMember Call ;

```

PROC GA Statement

```

PROC GA options ;

```

The PROC GA statement invokes the GA procedure. The following options are used with the PROC GA statement.

DATA n =SAS-data-set

specifies a data set containing data required to specify the problem, where n is an integer from 1 to 5. The data set is read and variables created matching the variables of the data set. If the data set has more than one observation, then the newly created variables are vector arrays with the size equal to the number of observations.

FIRSTGEN=SAS-data-set

specifies a SAS data set containing the initial solution generation. Different segments in the solution should be identified by variable names consisting of a letter followed by numbers representing the elements in the segments, in alphabetical order. For example, if the first segment of the solution uses real encoding and contains 10 elements, it should be represented by the numeric variables A1, A2, ..., A10. A second segment with integer encoding and five elements would be specified in the variables B1, B2, ..., B5. For Boolean encoding each Boolean element is represented by one variable in the data set, and for sequence encoding each position in the sequence is represented by one variable. If the data set contains a field named OBJECTIVE, then the value of that field becomes the objective value for that solution at initialization time (overriding the value computed by the input objective function), unless the field value is a missing value. The FIRSTGEN= and LASTGEN= options are designed to work together, so that a data set generated from a run of the GA procedure with a LASTGEN= option set can be specified in the FIRSTGEN= option of a subsequent run of the GA procedure, to continue the optimization from where it finished. If the number of observations in the data set is less than the population size specified in the [Initialize call](#), additional members are generated as specified in the [Initialize call](#) to complete the population. This feature makes it easy to seed an initial randomly generated population with chosen superior solutions generated by heuristics or a previous run of the GA procedure. If the data set contains more observations than the population size, the population is filled starting at the first observation, and the additional observations are not used.

LASTGEN=SAS-data-set

specifies a SAS data set into which the final solution generation is written. Different segments in the solution are identified by variable names consisting of a letter followed by numbers representing the elements in the segments, in alphabetical order. For example, if the first segment of the solution uses real encoding and contains 10 elements, it would be represented by the numeric variables A1, A2, ..., A10. A second segment with integer encoding and five elements would be specified in the variables B1, B2, ..., B5. For Boolean encoding each Boolean element is represented by one variable in the data set, and for sequence encoding each position in the sequence is represented by one variable. In addition to the solutions elements, the final objective value for each solution is output in the OBJECTIVE variable. The FIRSTGEN= and LASTGEN= options are designed to work together, so that a data set generated with a LASTGEN= option can be specified in the FIRSTGEN= option of a later run of the GA procedure.

LIBRARY=library-list

specifies a library or group of libraries for the procedure to search to resolve subroutine or function calls. The libraries can be created by using PROC FCMP and PROC PROTO. This option supplements the action of the CMPLIB= SAS option, but it permits you to designate libraries specific to this procedure invocation. Use the *libref.catalog* format to specify the two-level name of the library; *library-list* can be either a single library, a range of library names, or a list of libraries. The following examples demonstrate the use of the LIBRARY= option.

```
proc ga library = sasuser.xlib;
proc ga library = xlib1-xlib5;
proc ga library = (sasuser.xlib xlib1-xlib5 work.example);
```

MATRIX n =SAS-data-set

specifies a data set containing two-dimensional matrix data, where n is an integer from 1 to 5. A two-dimensional numeric array with the same name as the option is created and initialized from the data set. This option is provided to facilitate the input of tabular data to be used in setting up the optimization problem. Examples of data that might be provided by this option include a distance matrix for a traveling salesman problem or a matrix of coefficients for linear constraints.

MAXITER= n

specifies the maximum number of iterations to permit for the optimization process. A [ContinueFor](#) call overrides a limit set by this option.

NOVALIDATE= n

controls the amount of solution validity checking performed by the procedure. By default, the procedure verifies that valid solutions are being supplied at initialization and when the solution population is being updated by genetic operators or other user actions. If a solution segment has elements that exceed bounds set by a [SetBounds](#) call, those elements will be reset to the bound and a warning will be issued. If a solution segment contains other illegal values, an error will be signaled. This action is useful for maintaining the validity of the optimization and avoiding some errors that are often difficult to trace. However, this activity does consume CPU time, and you might want to use a strategy where you generate infeasible solutions at initialization or via genetic operators and then repair the solutions later in an update or objective routine. The NOVALIDATE= option enables you to do so, by turning off the validation checks. If n is 1, validation is turned off for initialization only, and if n is 2, validation is turned off for update only. If n is 3, all solution validity checking is turned off.

NOVALIDATEWARNING= n

controls the output of warning messages related to solution validation checking. Depending on the value of the NOVALIDATE= option, warning messages will be issued when the procedure repairs initial or updated solution segments to fit within bounds set with the [SetBounds](#) call. If n is 1, validation warnings are turned off for initialization only; and if n is 2, validation warnings are turned off for update only. If n is 3, all solution validation warnings are turned off.

SEED= n

specifies an initial seed to begin random number generation. This option is provided for reproducibility of results. If it is not specified, or if it is set to 0, a seed is chosen based on the system clock. The SEED value should be a nonnegative integer less than $2^{31} - 1$.

ContinueFor Call

```
call ContinueFor ( niter );
```

The ContinueFor call sets the number of additional iterations for the genetic algorithm optimization. The input to the ContinueFor subroutine is as follows:

<i>niter</i>	specifies that the optimization continue for <i>niter</i> more iterations. To stop at the current iteration, set <i>niter</i> to 0.
--------------	--

Cross Call

call Cross (*selected*, *seg*, *type*<, *parameter1*, *parameter2*, ... >) ;

The Cross call executes a genetic crossover operator from within a user subroutine. The inputs to the subroutine are as follows:

- selected* is an array that specifies the solutions to be crossed.
- seg* is the desired segment of the solution to which the crossover operator should be applied.
- type* is the type of crossover operator to apply, which also determines the number and type of parameters expected.
- parameter1-n* are optional parameters applicable to some operators.

The accepted values for *type* and the corresponding parameters are summarized in [Table 3.2](#).

Table 3.2 Crossover Operator Types

Type	Encodings	Parameters
'arithmetic'	real, integer	
'cycle'	sequence	
'heuristic'	real	
'null'	all encodings	
'order'	sequence	
'pmatch'	sequence	
'simple'	real, integer, Boolean	<i>alpha</i>
'twopoint'	real, integer, Boolean	<i>alpha</i>
'uniform'	real, integer, Boolean	<i>alpha</i> , <i>p</i>

The parameters are as follows:

- alpha* is a number such that $0 < \alpha \leq 1$.
- p* is a probability such that $0 < p \leq 0.5$.

The Cross call should be made only from within a user crossover subroutine. The precise action of these crossover operators is described in the section “[Crossover Operators](#)” on page 55.

Dynamic_array Call

call Dynamic_array (*arrayname*, *dim1*<, *dim2*, ..., *dim6*>) ;

The Dynamic_array call allocates a numeric array. The inputs to the Dynamic_array call are as follows:

- arrayname* is a previously declared array, whose dimensions are to be re-allocated.
- dim1* is the size of the first dimension.
- dim2*, ..., *dim6* are optional. Up to six dimensions can be specified.

The `Dynamic_array` call is normally used to allocate working arrays when the required size of the array is data-dependent. It is often useful in user routines for genetic operators or objective functions to avoid hard-coding array dimensions that might depend on segment length or population size. The array to be allocated must first be declared in an `ARRAY` statement with the expected number of dimensions, as in the following example:

```
subroutine sub(nx, ny);
  array x[1] /nosym;
  call dynamic_array(x, nx);
  array xy[1,1] /nosym;
  call dynamic_array(xy, nx, ny);
  ...
```

EvaluateLC Call

call EvaluateLC (*lc, results, sum, selected, seg*<, *child*>);

The EvaluateLC call evaluates linear constraints. The inputs to the EvaluateLC subroutine are as follows:

<i>lc</i>	is a two-dimensional array representing the linear constraints.
<i>results</i>	is a numeric array to receive the magnitude of the constraint violation for each linear constraint.
<i>sum</i>	is a variable to receive the sum of the constraint violations over all the constraints.
<i>selected</i>	is an array identifying the selected solution.
<i>seg</i>	is the segment of the solution to which the linear constraints apply.
<i>child</i>	is an optional parameter, and should be specified only when EvaluateLC is called from a user crossover operator.

The EvaluateLC routine can be called from a user crossover operator, mutation operator, or objective function to determine if a solution violates linear inequality constraints of the form $Ax \leq b$. For n linear constraints in m variables, the *lc* array should have dimension $n \times (m + 1)$. For each linear constraint $i = 1, \dots, n$, $lc[i, j] = A[i, j]$ for $j = 1, \dots, m$, and $lc[i, m + 1] = b[i]$. The *results* array should be one-dimensional with size n . The EvaluateLC call fills in the elements of *results* such that

$$results[i] = \begin{cases} 0, & \text{if } \sum_{j=1}^m A[i, j]x[j] \leq b[i] \\ \sum_{j=1}^m A[i, j]x[j] - b[i], & \text{otherwise} \end{cases}$$

In the variable *sum*, the EvaluateLC call returns the value $\sum_{i=1}^n results[i]$. Note that $sum \geq 0$, and $sum = 0$ implies that no constraints are violated. When you call EvaluateLC from your user routine, the *selected* parameter of the EvaluateLC call must be the same as the first parameter passed to your user routine to properly identify the solution to be checked. The *seg* parameter identifies which segment of the solution

should be checked. Real, integer, or Boolean encodings can be checked with this routine. If EvaluateLC is called from a user crossover operator, the *child* parameter must be specified to indicate which offspring is to be checked. The value *child* = 1 requests the first offspring, *child* = 2 requests the second, and so on.

GetDimensions Call

call GetDimensions (*source*, *dest*) ;

The GetDimensions call gets the dimensions of an array variable. The inputs to the GetDimensions subroutine are as follows:

source is the array variable whose dimensions are desired.

dest is an array to receive the dimensions of *source*.

The GetDimensions subroutine is used to get the dimensions of an array passed into a user subroutine. The input *dest* should be an array of one dimension, with at least as many elements as there are dimensions in *source* (a maximum of 6). Any extra elements in *dest* are filled with zeros.

GetObjValues Call

call GetObjValues (*dest*, *n*) ;

The GetObjValues call retrieves objective function values from the current solution generation. The inputs to the GetObjValues subroutine are as follows:

dest is an array to receive the objective values.

n is the number of objective values to get. The GetObjValues subroutine is used to retrieve the objective values for the current solution generation. It can be called from a user update routine or finalize routine. If it is called from a finalize routine, and if the *elite* parameter from a SetElite call is 1 or greater, then the first *elite* members of the population are the fittest of the population, and they are sorted in order, starting with the most fit. The input *dest* should be a dimensioned variable, with dimension greater than or equal to *n*.

GetSolutions Call

call GetSolutions (*sol*, *n*, *seg*) ;

The GetSolutions call retrieves solutions from the current generation. The inputs to the GetSolutions subroutine are as follows:

sol is an array to receive the solution elements.

n is the number of solutions to get.

seg is the segment of the solution to retrieve.

The GetSolutions subroutine is used to retrieve solutions from the current generation. You would normally call it from an update or finalize subroutine for post processing or analysis. If the *elite* parameter has been set with a SetElite call, then the first *elite* members of the population are the fittest, and they are sorted in order, starting with the most fit. The *sol* variable should have two dimensions, with the first dimension

representing the solution number, and the second representing the element within the solution. For example, if the encoding of the problem was I10, then *sol*[2, 3] would be the value of the third element of the second solution in the current population. For real, integer, Boolean, and sequence encoding, each solution element is mapped to the corresponding element of the *sol* array. The *seg* parameter specifies the solution segment desired. For example, if the encoding was set in the [SetEncoding call](#) to 'R10I5', then segment 1 is R10 and segment 2 is I5.

Initialize Call

call Initialize (*option*, *size* < ,*option*, *size* > ...) ;

The Initialize call creates the initial solution generation. The inputs to the Initialize subroutine are as follows:

<i>option</i>	is a string that specifies an initialization option.	The Initialize subroutine must be called to create the first solution generation, and can be used to reinitialize a solution population during the optimization process. The available options and their effect are as follows:
<i>size</i>	is the number of solutions to create by using a given option.	

- | | |
|----------------|---|
| ‘_uniform_’ | generate uniformly distributed solutions—for integer and real encoded segments for which SetBounds has been called, segment elements will be uniformly distributed between the upper and lower bounds. If no bounds have been specified for an integer or real encoded segment, then the elements will be set to 0. For Boolean encoded segments the elements will be randomly assigned 0 or 1, and for sequence encoded segments random sequences will be generated. |
| ‘_dataset_’ | read solutions from the data set specified in a FIRSTGEN= option. If the data set has more observations than requested, the extra observations are ignored. |
| ‘default’ | read solutions from the data set specified in a FIRSTGEN= option, if one was specified. If none was specified or the data set has fewer observations than requested, fill in the remaining solution population by using the ‘_uniform_’ option. |
| ‘_retain_’ | bring forward the best solutions from the current generation. This option cannot be used for the first initialization. |
| ‘user-routine’ | Any string not matching the preceding options is interpreted to be a user-defined initialization routine. See the section “ Defining a User Initialization Routine ” on page 69 for information about defining an initialization subroutine. |

After the Initialize call, the current solution population size is the sum of the population sizes specified for each option. The following rules also apply to the option specifications:

1. All options must be literal quoted strings.
2. No option type can be specified more than once in the same Initialize call. No more than one user initialization routine can be specified.

3. 'default' cannot be specified in combination with the '_uniform_' or '_dataset_' options.
4. If the '_uniform_' option is specified, the solution encoding must include at least one segment that is either Boolean or sequential, or that has bounds specified with a [SetBounds](#) call.

MarkPareto Call

call MarkPareto (*result, n, objectives, minmax*) ;

The MarkPareto call identifies the Pareto-optimal set from a population of solutions. The inputs to the MarkPareto call are as follows:

<i>result</i>	is a one-dimensional array to accept the results of the evaluation. Its size should be the same as the size of the population being evaluated; $result[i] = 1$ if solution i is Pareto optimal, and 0 otherwise.
<i>n</i>	is a variable to receive the number of Pareto-optimal solutions.
<i>objectives</i>	is a two-dimensional array that contains the multiple objective values for each solution. It should be dimensioned $[p, q]$, where p is the size of the population, and q is greater than or equal to the number of objectives to be considered.
<i>minmax</i>	is a one-dimensional array to specify how the objective values are to be used. It should be of size q , where q is greater than or equal to the number of objectives to be considered. $minmax[k] = -1$ if objective k is to be minimized, $minmax[k] = 1$ if objective k is to be maximized, $minmax[k] = 0$ if objective k is not to be considered, and $minmax[k] = -2$ designates an objective that prevents the member from being considered for Pareto optimality if it is nonzero.

The MarkPareto call is used to identify the Pareto-optimal subset from a population of solutions. See the section “[Optimizing Multiple Objectives](#)” on page 73 for a full discussion of Pareto optimality. MarkPareto can be called from a user update routine, which is called after the individual solution objective values have been calculated and before selection takes place. To make best use of this routine, in your encoding you need to set up a segment to record all the objective values you intend to use in the Pareto-optimal set evaluation. In a user objective function, you should calculate the multiple objectives and write them to the chosen segment. In an update routine (designated with a [SetUpUpdateRoutine](#) call), you can use a [GetSolutions](#) call to retrieve these segments, and then pass them to a MarkPareto call. The following statements shows how this could be done:

```
subroutine update(popsize);

    array objectives[1,1] /nosym;
    call dynamic_array(objectives, popsize, 3);

    array pareto[1] /nosym;
    call dynamic_array(pareto, popsize);

    array minmax[3] /nosym (1 -1 0);
```

```

call GetSolutions(objectives, popsize, 2);

call MarkPareto(pareto, npareto, objectives, minmax);

do i = 1 to popsize;
    objectives[i,3] = pareto[i];
end;

call UpdateSolutions(objectives, popsize, 2);

call SetElite(npareto);

endsub;

```

This is an example of a user update routine that might be used in a multiobjective optimization problem. It is assumed that a user objective function has calculated two different objectives and placed their values in the first two elements of segment 2 of the problem encoding. The first objective is to be maximized, and the second is to be minimized. Segment 2 has three elements, and the third element is used to mark the Pareto-optimal solutions. After dynamically allocating the necessary arrays from the *popsize* (population size) parameter, the update routine first retrieves the current solutions into the *objectives* array with the [GetSolutions call](#). It then passes the *objectives* array directly to the MarkPareto call to perform the Pareto-optimal evaluations. Note that the *minmax* array directs the MarkPareto call to maximize the first element, minimize the second, and ignore the third. After the MarkPareto call, the update routine writes the results back to the third element of the *objectives* array, and writes the *objectives* array back to the solution population with the UpdateSolutions call. This marks the solutions that compose the Pareto-optimal set. The update routine then sets the *elite* parameter equal to the number of Pareto-optimal solutions with the [SetElite call](#). It is assumed that the user has provided a fitness comparison function (designated with a [SetCompareRoutine call](#)) that always selects a Pareto-optimal solution over a non-Pareto-optimal one, so the *elite* setting guarantees that all the Pareto-optimal solutions are retained from generation to generation. [Example 3.3](#) illustrates the use of the MarkPareto call.

Mutate Call

call Mutate (*selected*, *seg*, *type*<, *parameter1*, *parameter2*, ... >);

The Mutate call executes a genetic mutation operator from within a user subroutine. The inputs to the subroutine are as follows:

- selected* is an array that specifies the solution to be mutated.
- seg* is the desired segment of the solution to which the mutation should be applied.
- type* is the type of mutation operator to apply, which also determines the number and type of parameters expected.
- parameter1-n* are optional parameters applicable to some operators.

The accepted values for *type* and the corresponding parameters are summarized in [Table 3.3](#).

Table 3.3 Mutation Operator Types

Type	Encodings	Parameters
'delta'	real, integer	<i>delta</i> <i>n</i>
'invert'	sequence	
'swap'	sequence	<i>n</i>
'uniform'	real, integer, Boolean	<i>np</i>

The parameters are as follows:

<i>delta</i>	is a vector of delta values for each component of the solution, used only for the Delta mutation operator.
<i>n</i>	specifies the number of solution elements that should be mutated for the Delta operator, and the number of swaps that should be made for the Swap operator.
<i>np</i>	specifies the number of solution elements that should be mutated, if <i>np</i> is integer; specifies the mutation probability for each solution element if $0 < np < 1$.

The Mutate call should be made only from within a user mutation subroutine. The precise action of these mutation operators is described in the section “[Mutation Operators](#)” on page 61.

Objective Function

r = Objective (*selected*, *seg*, *type*<, *parameter1*, *parameter2*, ... >) ;

The Objective call evaluates a standard objective function from within a user subroutine. The inputs to the function are as follows:

<i>selected</i>	is an array that specifies the solution to be evaluated.
<i>seg</i>	is the desired segment of the solution to be evaluated.
<i>type</i>	is objective function name, which also determines the number and type of parameters expected.
<i>parameter1-n</i>	are optional parameters applicable to particular objective functions.

The accepted values for *type* and the corresponding parameters are summarized in [Table 3.4](#).

Table 3.4 Objective Function Types

type	encodings	parameters
'TSP'	sequence	<i>distances</i>

The parameters are as follows:

<i>distances</i>	is a matrix of distances between locations, such that <i>distances</i> [<i>i</i> , <i>j</i>] is the distance between location <i>i</i> and <i>j</i> .
------------------	---

The Objective call should be made only from within a user objective function. The precise actions of the standard objective functions are described in the section “[Objective Functions](#)” on page 63.

PackBits Call

call PackBits (*array*, *start*, *width*, *value*);

The PackBits call writes bits to a packed integer array. The inputs to the PackBits subroutine are as follows:

<i>array</i>	is an array to which the value is to be assigned.
<i>start</i>	is the starting position for the bit assignments.
<i>width</i>	is the number of bits to assign.
<i>value</i>	is the value to be assigned to the bits. For a single bit, this should be 0 or 1.

The PackBits subroutine facilitates the assignment of bit values into an integer array, effectively using the integer array as an array of bits. One common use for this routine is within a user genetic operator subroutine to pack bit values into an integer vector solution segment. The bit values assigned with the PackBits call can be retrieved with the [UnpackBits](#) function.

The *start* parameter is the lowest desired bit position in the bit vector, where the least significant bit of *value* is to be stored. The *start* parameter can range in value from 1 to *maxbits*, where *maxbits* is the product of 32 times the number of elements in the integer array.

The *width* parameter is the number of bits of *value* to be stored. It is bounded by $0 < width < (maxbits - start + 1)$ and must also not exceed 32.

Bits not within the range defined by *start* and *width* are not changed. If the magnitude of *value* is too large to express in *width* bits, then the *width* least significant bits of *value* are transferred. The following program fragment, which might occur in a mutation subroutine, first reads a selected solution segment into *s* with the [ReadMember](#) call and then overwrites the first and second least significant bits of the solution with ones before writing it back to the current generation.

```
array s[2];
call ReadMember(selected, seg, s);
...
/* intervening code */
...
call PackBits(s, 1, 2, 3); /* start = 1, width = 2,
                           * value = 3 = binary 11
                           */
call WriteMember(selected, seg, s);
```

Programming Statements

This section lists the programming statements used to initialize the model, code the objective function, and control the optimization process in the GA procedure. It documents the differences between programming

statements in the GA procedure and programming statements in the DATA step. The syntax of programming statements used in **PROC GA** is identical to that of programming statements used in the FCMP procedure.

Most of the programming statements that can be used in the SAS DATA step can also be used in the GA procedure. See the *SAS Statements: Reference* or BASE SAS documentation for a description of the SAS programming statements.

```

variable = expression;
variable + expression;
arrayvar[subscript] = expression;
ABORT;
CALL subroutine < ( parameter-1 <, ...parameter-n > ) >;
DELETE;
DO program-statements; END;
DO variable = expression TO expression <BY expression>;
    program-statements; END;
DO WHILE expression ;
    program-statements; END;
DO UNTIL expression ;
    program-statements; END;
GOTO statement_label ;
IF expression THEN program-statement;
    <ELSE program-statement>;
PUT < variable(s)> <@ | @@> ;
RETURN <(expression)>;
SELECT <(select-expression)>;
    WHEN-1 (expression-1 <...expression-n>)program-statement ;
    <WHEN-n (expression-1 <...expression-n>)program-statement ;>
    <OTHERWISE program-statement ;>
STOP;
SUBSTR( variable, index, length ) = expression;

```

For the most part, the SAS programming statements work as they do in the SAS DATA step as documented in the *SAS Statements: Reference*. However, there are several differences that should be noted.

- The ABORT statement does not permit any arguments.
- The DO statement does not permit a character index variable. Thus

```
do i = 1,2,3;
```

is supported; however,

```
do i = 'A', 'B', 'C' ;
```

is not.

- The PUT statement, used mostly for program debugging in **PROC GA**, supports only some of the features of the DATA step PUT statement, and has some new features that the DATA step PUT statement does not:
 - The PROC GA PUT statement does not support line pointers, factored lists, iteration factors, overprinting, _INFILE_, the colon (:) format modifier, or “\$”.

- The PROC GA PUT statement does support expressions, but the expression must be enclosed inside parentheses. For example, the following statement displays the square root of *x*:

```
put (sqrt(x));
```
- The PROC GA PUT statement permits an array name without subscripts. The statement `PUT A;` prints all the elements of array *A*. The statement `PUT A=;` prints the elements of array *A* with each value labeled with the name of the element variable.
- The PROC GA PUT statement supports the print item `_PDV_` to print a formatted listing of all variables in the program. For example, the following statement displays a more readable listing of the variables than the `_all_` print item:

```
put _pdv_;
```
- The WHEN and OTHERWISE statements permit more than one target statement. That is, DO/END groups are not necessary for multiple-statement WHENs. For example, the following syntax is valid:

```
SELECT;
WHEN ( exp1 )  stmt1;
                stmt2;
WHEN ( exp2 )  stmt3;
                stmt4;
END;
```

ReadChild Call

call ReadChild (*selected*, *seg*, *n*, *values*);

The ReadChild call reads a segment from a selected child solution into an array, within a user crossover operator. The inputs to the ReadChild subroutine are as follows:

<i>selected</i>	specifies the family (parents and children) obtained from the selection process.
<i>seg</i>	specifies the solution segment to be read.
<i>n</i>	specifies the child in the family from which to read the solution segment.
<i>values</i>	specifies an array to receive the solution elements.

The ReadChild call is used to obtain the solution values for manipulation within a user crossover operator subroutine. Normally it is needed only if you need to augment the action of a GA procedure-supplied crossover operator. You might need to make modifications to satisfy constraints, for example. The *selected* parameter is passed into the user subroutine by the GA procedure. The *seg* parameter is the desired segment of the solution to be obtained. Segments, which correspond to different encodings in the encoding string, are numbered, starting from 1 as the first segment. The parameter *n* should be 1 to get the first child and 2 for the second. The parameter *values* is an array, which should be dimensioned large enough to contain the segment's encoding. For example, the following subroutine illustrates how you could use the Read/WriteChild calls to modify offspring generated with a standard genetic operator:

```
call SetEncoding('R5');

subroutine cross(selected[*]);
```

```

/* generate offspring with arithmetic crossover operator */
call CrossArithmetic(selected, 1); /* here 1 refers to segment 1*/

array child1[5];
array child2[5];

/* get elements of first child solution */
call ReadChild(selected, 1, 1, child1);

/* get elements of second child solution values */
call ReadChild(selected, 1, 2, child2);

...
/* code to modify elements in child1 and child2 */
...

call WriteChild(selected, 1, 1, child1);
call WriteChild(selected, 1, 2, child2);

```

ReadCompare Call

call ReadCompare (*selected, seg, n, values*) ;

The ReadCompare call reads a segment from a selected solution into an array, within a user fitness comparison subroutine. The inputs to the ReadCompare subroutine are as follows:

<i>selected</i>	specifies the pair of solutions to be compared, obtained from the selection process.
<i>seg</i>	specifies the solution segment to be read.
<i>n</i>	specifies the solution (1 or 2) from which to read the segment.
<i>values</i>	specifies an array to receive the solution elements.

The ReadCompare call is used to obtain the solution values for manipulation within a user fitness comparison subroutine, which can be designated in a [SetCompareRoutine call](#).

ReadMember Call

call ReadMember (*selected, seg, destination*) ;

The ReadMember call reads the selected solution into an array for a user objective function or mutation operator.

The inputs to the ReadMember subroutine are as follows:

<i>selected</i>	is a parameter passed to the user subroutine by the GA procedure, which points to the selected solution.
<i>seg</i>	specifies which segment of the solution to retrieve.
<i>destination</i>	specifies an array in which to store the solution elements.

The ReadMember call is used within a user objective function or mutation operator to obtain the elements of a selected solution and write them into a specified vector. They can then be used to compute an objective value, or in the case of a mutation operator, manipulated and written back out with a [WriteMember call](#).

ReadParent Call

call ReadParent (*selected, seg, n, destination*) ;

The ReadParent call reads selected solution elements into an array in a user crossover subroutine. The inputs to the ReadParent subroutine are as follows:

- selected* is a parameter passed to the user subroutine by the GA procedure, which points to the selected solution family.
- seg* is the segment of the desired parent solution to be obtained.
- n* is the number of the parent, starting at 1.
- destination* is an array in which to store the solution elements.

The ReadParent subroutine is called inside a user crossover operator subroutine to obtain the elements of selected parent solutions. Normally you would then manipulate and combine the elements of the two parents and use a [WriteChild call](#) to create the child offspring and complete the action of the crossover operator.

ReEvaluate Call

call ReEvaluate (<*index*>) ;

The ReEvaluate call reruns the evaluation phase of the genetic algorithm. The inputs to the ReEvaluate subroutine are as follows:

- index* is a numeric scalar or array that specifies the indices of the solutions to be updated. The indices correspond to the order of the solutions obtained from a GetSolutions call.

The ReEvaluate call recomputes the objective values for the current generation. You do not normally need to use this call, because the GA procedure evaluates the objective function during the optimization process in the evaluation phase. This subroutine should be called from a user update or finalize routine if a parameter that affects the objective value or solution is changed. The optional *index* parameter enables you to restrict the recomputation to the solution or subset of solutions specified. If the *index* parameter is not supplied, then the objective values of all the solutions will be recomputed.

For example, you might have a user objective function that can perform an additional local optimization if a particular parameter is set. If your update routine changes that parameter, then you should call the ReEvaluate subroutine to update the solutions and objective function values.

SetBounds Call

call SetBounds (*lower, upper* <, *seg*>) ;

The SetBounds call sets constant upper and lower bounds. The inputs to the SetBounds subroutine are as follows:

- lower* is a lower bound for the solution components.
- upper* is an upper bound for the solution components.
- seg* is optional, and specifies a segment of the solution to which the bounds apply. If *seg* is not specified, then it defaults to a value of 1.

The SetBounds subroutine is used to establish upper and lower bounds on the solution space. It applies only to integer and real encoding. For multiple segment encoding, use the *seg* parameter to specify a segment other than the first. *upper* and *lower* must be arrays, with the same dimension as the encoding size. SetBounds must be called for all integer or real encoded segments to which you apply the Uniform mutation operator. The action of the standard mutation and crossover operators supplied by the GA procedure is automatically modified so that the bounds established by a SetBounds call are respected. For an integer encoded segment, only integer values are allowed for the upper and lower bounds.

SetCompareRoutine Call

call SetCompareRoutine ('routine') ;

The SetCompareRoutine call installs a user function to compare the fitness of solutions. The input to the SetCompareRoutine subroutine is as follows:

- routine* is the name of a function you have defined, which is called when necessary to compare the fitness of solutions. This parameter must be a string literal; a variable is not accepted.

The SetCompareRoutine call enables you to designate a function you have defined to be used in the selection process when comparing the fitness of two solutions. The selector options that involve ranking solutions by fitness, including tournament and duel selection, will use the designated function instead of comparing the objective values directly. If the SetCompareRoutine is not called, or if it is called with an input value of 'default', then the objective function value will be used to compare solution fitness. The SetCompareRoutine call provides you with a way to factor multiple fitness criteria into the selection process. See the section “[Defining a User Fitness Comparison Routine](#)” on page 63 for a full description of how the fitness comparison routine should be structured and what it should return. You can use this feature to implement [multiobjective optimization](#) and other advanced optimization strategies.

SetCross Call

call SetCross (type< , seg>< , pname, pvalue>< , pname, pvalue>...) ;

The SetCross call sets the crossover operator. The inputs to the SetCross subroutine are as follows:

<i>type</i>	is the name of the crossover operator to be applied.
<i>seg</i>	is optional, and specifies a segment of the solution to which the operator should be applied. If <i>seg</i> is not specified, then it defaults to a value of 1. <i>seg</i> needs to be specified only if multisegment encoding is used.
<i>pname</i>	is optional, and specifies the name of a particular property to be set for the crossover operator.
<i>pvalue</i>	specifies the value to be assigned to the corresponding property name.

The SetCross routine is used to assign a standard crossover operator. For multisegment encoding, the operator can be assigned to a particular solution segment with the *seg* parameter; otherwise a default segment of 1 is assumed. You can set different crossover operators for different segments with multiple SetCross calls. When a crossover event occurs, all segments in the parent solutions for which a crossover operator has been designated will undergo crossover. If more than one SetCross call is made for the same segment, then the last call nullifies any previous call for that segment. Also, a SetCrossRoutine call nullifies all previous SetCross calls, and a SetCross call nullifies a previous SetCrossRoutine call. Properties for the chosen crossover operator can be set with optional *pname-pvalue* pairs. It is also possible to set or reset operator properties with a SetProperty call.

The accepted values for *type* and the corresponding properties are summarized in Table 3.5. See the section “Crossover Operators” on page 55 for a full description of the available operators.

Table 3.5 Crossover Operator Types and Properties

type	encodings	properties
‘arithmetic’	real, integer	
‘cycle’	sequence	
‘heuristic’	real	
‘order’	sequence	
‘pmatch’	sequence	
‘simple’	real, integer, Boolean	‘alpha’
‘twopoint’	real, integer, Boolean	‘alpha’
‘uniform’	real, integer, Boolean	‘alpha’, ‘p’

SetCrossProb Call

call SetCrossProb (*p*) ;

The SetCrossProb call sets the crossover probability. The input to the SetCrossProb subroutine is as follows:

p is the crossover probability.

The SetCrossProb subroutine is used to set the crossover probability for the genetic algorithm optimization process. The crossover probability *p* should be between 0 and 1. Typical values for this parameter range from 0.6 to 1.0. The crossover probability will be overridden if required by a SetElite call. The elite solutions are passed on to the next generation without undergoing crossover, regardless of the crossover probability.

SetCrossRoutine Call

call SetCrossRoutine (*'routine'*< , *nparents*, *nchildren*>) ;

The SetCrossRoutine call installs a user subroutine for the crossover operator. The inputs to the SetCrossRoutine subroutine are as follows:

<i>routine</i>	is the name of a subroutine you have defined, which is called when the mutation operator is applied. This parameter must be a string literal; a variable is not accepted.
<i>nparents</i>	is optional, and specifies the number of parent solutions the operator requires. If not specified, 2 is assumed.
<i>nchildren</i>	is optional, and specifies the number of children solutions the operator will generate. If not specified, 2 is assumed.

SetElite Call

call SetElite (*elite*) ;

The SetElite call sets the number of best solutions to pass to the next generation. The input to the SetElite subroutine is as follows:

<i>elite</i>	is the number of best solutions to be passed unmodified from the current solution generation to the next.
--------------	---

The SetElite subroutine is used to ensure that the best solutions encountered in the optimization are not lost by the random selection process. In pure tournament selection, although better solutions are more likely to be selected, it is also possible that any given solution will not be chosen to participate in a tournament, and even if it is selected, it might be modified by crossover or mutation. The SetElite call modifies the optimization process such that the best *elite* solutions in the current population are exactly preserved and passed on to the next generation. This behavior is observed regardless of the crossover or mutation settings. When a SetElite call is made, the first *elite* solutions in the population retrieved by a [GetSolutions call](#) or output to a data set are the fittest, and these *elite* solutions are sorted so that the most fit is first. In general, using the SetElite call speeds the convergence of the optimization process. However, it can also lead to premature convergence before a true global optimum is reached. If no SetElite call is made, a default *elite* value of 1 is used by the GA procedure to make sure that the best solution encountered in the optimization process is never lost.

SetEncoding Call

call SetEncoding (*encoding*) ;

The SetEncoding call specifies the problem encoding. The input to the SetEncoding subroutine is as follows:

<i>encoding</i>	is a string used to specify the form of the solution.
-----------------	---

The SetEncoding subroutine is used to establish the type of problem solution encoding. The *encoding* parameter should be a string of letter-number pairs, where the letter determines the type of encoding: I for integer, R for real-valued, S for sequences, and B for Boolean values. Each letter is followed by a

number to indicate the number of components for that encoding. Multiple letter-number pairs can be used to specify a multisegment encoding. For example, the following call specifies that solutions be in the form of a 10-member integer vector:

```
call SetEncoding('I10');
```

The following call specifies that solutions have a 5-component integer segment and a 10-component real-valued segment:

```
call SetEncoding('I5R10');
```

See the section “[Using Multisegment Encoding](#)” on page 53 for details about using multisegment encoding.

SetFinalize Call

```
call SetFinalize ( 'routine' );
```

The SetFinalize call designates a user subroutine to perform post processing at the end of the optimization process. The input to the SetFinalize subroutine is as follows:

<i>routine</i>	is the name of a subroutine you have defined, which is called when the optimization process ends. This parameter must be a string literal; a variable is not accepted.
----------------	--

The SetFinalize subroutine enables you to define a subroutine to be called at the end of the optimization process. You might use this subroutine to perform additional refinements of the best solution, or you could generate and write out additional data for plots or reports.

SetMut Call

```
call SetMut type<, seg><, pname, pvalue><, pname, pvalue>...);
```

The SetMut call sets the mutation operator. The inputs to the SetMut subroutine are as follows:

<i>type</i>	is the name of the mutation operator to be applied.
<i>seg</i>	is optional, and specifies a segment of the solution to which the operator should be applied. If <i>seg</i> is not specified, then it defaults to a value of 1. <i>seg</i> needs to be specified only if multisegment encoding is used.
<i>pname</i>	is optional, and specifies the name of a particular property to be set for the mutation operator.
<i>pvalue</i>	specifies the value to be assigned to the corresponding property name.

The SetMut routine is used to assign a standard mutation operator. For multisegment encoding, the operator can be assigned to a particular solution segment with the *seg* parameter; otherwise a default segment of

1 is assumed. You can set different mutation operators for different segments with multiple SetMut calls. When a mutation event occurs, all the operators will be applied to the same solution. If more than one SetMut call is made for the same segment, then the last call nullifies any previous call for that segment. Also, a SetMutRoutine call nullifies all previous SetMut calls, and a SetMut call nullifies a previous SetMutRoutine call. Properties for the chosen mutation operator can be set with optional *pname-pvalue* pairs. It is also possible to set or reset operator properties with a SetProperty call.

The accepted values for *type* and the corresponding properties are summarized in Table 3.6. See the section “Mutation Operators” on page 61 for a full description of the available operators.

Table 3.6 Mutation Operator Types and Properties

type	encodings	properties
'delta'	real, integer	'delta' 'nchange'
'invert'	sequence	
'null'	all encodings	
'swap'	sequence	'nswap'
'uniform'	real, integer, Boolean	'nchange' 'pchange'

SetMutProb Call

call SetMutProb (*p*) ;

The SetMutProb call sets the mutation probability. The input to the SetMutProb subroutine is as follows:

p is the mutation probability.

The SetMutProb subroutine is used to set the mutation probability for the genetic algorithm optimization. The probability *p* should be a number between 0 and 1, and is interpreted as the probability that a solution in the next generation should have the mutation operator applied to it. If a SetElite call has been made, then the elite solutions do not undergo mutation. Generally, a high mutation probability degrades the convergence of the genetic algorithm optimization, but some level of mutation is required to assure a thorough search and avoid premature convergence before the global optimum is found. Typical values for *p* are near 0.05 or less.

SetMutRoutine Call

call SetMutRoutine ('*routine*') ;

The SetMutRoutine call installs a user subroutine for the mutation operator. The input to the SetMutRoutine subroutine is as follows:

routine is the name of a subroutine you have defined, which is called when the mutation operator is applied. This parameter must be a string literal; a variable is not accepted.

The SetMutRoutine call enables you to designate a subroutine you have defined to be used for the mutation operator. Your subroutine will be called whenever the mutation operation is performed. See the section “[Defining User Genetic Operators](#)” on page 64 for more information about defining a mutation operator.

SetObj Call

call SetObj (*type*, *minmax*, *<*, *seg*>*<*, *pname*, *pvalue*>*<*, *pname*, *pvalue*>...) ;

The SetObj call sets the objective function. The inputs to the SetObj subroutine are as follows:

<i>type</i>	is the name of the objective function to be used.
<i>minmax</i>	is an indicator to maximize or minimize the objective. A value of 0 is used to specify a minimization, and a value of 1 to specify maximizing the objective.
<i>seg</i>	is optional, and specifies a segment of the solution to which the objective function should be applied. If <i>seg</i> is not specified, then it defaults to a value of 1. <i>seg</i> needs to be specified only if multisegment encoding is used.
<i>pname</i>	is optional, and specifies the name of a particular property to be set for the objective function.
<i>pvalue</i>	specifies the value to be assigned to the corresponding property name.

The SetObj routine is

used to assign a procedure-supplied objective function. For multisegment encoding, the objective can be assigned to a particular solution segment with the *seg* parameter; otherwise a default segment of 1 is assumed. If more than one SetObj call is made, then the last call nullifies any previous call. Also, a [SetObjFunc](#) call nullifies all previous SetObj calls, and a SetObj call nullifies a previous [SetObjFunc](#) call. Properties for the chosen objective can be set with optional *pname-pvalue* pairs. It is also possible to set or reset objective properties with a [SetProperty](#) call.

The accepted values for *type* and the corresponding properties are summarized in [Table 3.7](#). See the section “[Objective Functions](#)” on page 63 for a full description of the available objectives.

Table 3.7 Objective Function Types and Properties

type	encodings	properties
'TSP'	real, integer	'distances'

SetObjFunc Call

call SetObjFunc ('*fname*', *minmax*) ;

The SetObjFunc call sets the objective to a user-defined function. The inputs to the SetObjFunc subroutine are as follows:

- fname* is the name of a user objective function. This parameter must be a literal string.
- minmax* is set to 0 to minimize the objective, 1 to maximize.

The SetObjFunc subroutine is used to designate a user function to be the objective for the optimization process. The SetObjFunc call accepts a literal string only for the function name; you cannot use a variable or expression. See the section “[Defining an Objective Function](#)” on page 68 for more information about defining your own objective function. If multiple SetObjFunc calls are made, only the last one is in effect; the last call nullifies any previous SetObjFunc or [SetObj](#) calls.

SetProperty Call

call SetProperty (*otype* < , *seg* > , *pname* , *pvalue* < , *pname* , *pvalue* > . . .) ;

The SetProperty call modifies properties of genetic operators, objective functions, and selectors. The inputs to the SetProperty subroutine are as follows:

- otype* is the type of operator. It should have a value of ‘cross’ for a crossover operator, ‘mut’ for a mutation operator, ‘obj’ for an objective function, or ‘sel’ for a selector.
- seg* is optional, used only for mutation and crossover operators, and specifies the segment in which the operator resides. It is necessary only for multisegment encoding. The default value if *seg* is not specified is 1.
- pname* specifies the name of a particular property to be set.
- pvalue* specifies the value to be assigned to the corresponding property name. Multiple property name-value pairs can be supplied in a SetProperty call.

The SetProperty call is used to set or modify properties of a genetic operator, objective function, or a selector. It can be called anytime during the optimization process to dynamically adapt optimization parameters. For example, you might call SetProperty from a user [update](#) routine to reduce the magnitude of the *delta* vector of a *delta* mutation operator as the optimization progresses to an optimum.

SetSel Call

call SetSel (*selector* < , *pname* , *pvalue* > < , *pname* , *pvalue* > . . .) ;

The SetSel call sets the selection parameters. The inputs to the SetSel subroutine are as follows:

- selector* is the type of selection strategy to be used.
- pname* is optional, and specifies the name of a particular property to be set for the selector operator.
- pvalue* specifies the value to be assigned to the corresponding property name.

The SetSel call is used to specify a selector for the regeneration process, which selects members of the current generation to be propagated to the next. Generally, selection is based on solution fitness, with the fittest solutions more likely to be selected.

The supported values for *selector* and the corresponding selector properties and their default values are summarized in Table 3.8. See the section “Specifying the Selection Strategy” on page 71 for a full description of the available selectors and their properties.

Table 3.8 Selectors and Properties

Selector	Properties	Default values
'tournament'	'size'	2
'duel'	'pbest'	0.8

SetUpdateRoutine Call

call SetUpdateRoutine ('routine');

The SetUpdateRoutine call designates a control subroutine to be called at each iteration. The input to the SetUpdateRoutine subroutine is as follows:

routine is the name of a subroutine you have defined that is called once during each iteration of the optimization process. This parameter must be a string literal; a variable is not accepted.

The SetUpdate subroutine enables you to define a subroutine to be called at each iteration of the optimization process, in order to monitor the progress of the genetic algorithm, adjust optimization parameters, or perform calculations that depend on the population as a whole. The specified routine is called once at each iteration, just before the selection process and after the evaluation phase. See the section “Defining a User Update Routine” on page 67 for a discussion of how an update routine might be used.

ShellSort Call

call ShellSort (*x*, <, *by*<, *descend*> >);

The ShellSort call sorts a numeric array. The inputs to the ShellSort subroutine are as follows:

x is a one or two dimensional array to be sorted.

by is an optional numeric scalar or array that specifies the columns by which the array is to be sorted. If not specified, column 1 is the default.

descend is an optional numeric scalar or array used to specify which columns in the *by* parameter are to be in descending order. Any columns not specified in a *descend* parameter will be in ascending order.

The ShellSort subroutine sorts the *x* array by the columns specified in the *by* parameter, with the first column having highest precedence, and subsequent columns applied within the preceding *by* groups. Sorting will be

done in ascending order for each column unless that column is also specified in the *descend* parameter. In general the ShellSort routine does not preserve the original order in case of ties. For example, the following statements sort an array *x* into ascending order with respect to the first column, and sort groups with the same first column value by descending order of third column values:

```
array x[100, 3] /nosym;

...

array by[2] /nosym;
by[1] = 1;
by[2] = 3;
descend = 3;
call ShellSort(x, by, descend);
```

Shuffle Call

call Shuffle (*x*);

The Shuffle call randomly reorders a numeric array. The input to the Shuffle subroutine is as follows:

x is a numeric array to be randomly shuffled.

The Shuffle subroutine randomly rearranges the elements of the *x* array. One example of where it might be used is in a user-supplied initialization routine, to generate a random sequence-encoded solution segment.

UnpackBits Function

r = UnpackBits (*source*, *start*, *width*);

The UnpackBits call retrieves bit values from a packed integer array. The inputs to the UnpackBits function are as follows:

source is an array containing the packed bit values.

start is the starting bit, with the lowest bit starting at 1.

width is the number of bits to retrieve. A value of 1 retrieves a single bit.

The UnpackBits function facilitates the extraction of bit values from arbitrary locations in an integer array. One common use for it is to retrieve bit values from an integer solution segment in a user objective or genetic operator routine. The return value, *start*, and *width* parameters are consistent with the [PackBits](#) call, which you can use to store bit values to an integer array.

The *start* parameter is the lowest desired bit position in the bit vector, corresponding to the least significant bit of the return value. The *start* parameter can range in value from 1 to *maxbits*, where *maxbits* is the product of 32 times the number of elements in the integer array.

The *width* parameter is the number of bits to be read into the return value from the array. It is bounded by $0 < width < (maxbits - start + 1)$, and must also not exceed 32.

UpdateSolutions Call

call UpdateSolutions (*sol*, *n*, *seg*) ;

The UpdateSolutions call updates the current solution population. The inputs to the UpdateSolutions subroutine are as follows:

- sol* is an array containing the replacement solution elements.
- n* is the number of solutions to update.
- seg* is the segment of the solution to replace.

The UpdateSolutions subroutine is used to replace the values of the selected solution segment with new values computed in an update routine. The update routine can be designated in a [SetUpUpdateRoutine](#) call. The UpdateSolutions call is often used to implement advanced strategies such as marking Pareto-optimal sets or employing local optimizations. The *sol* parameter should have 2 dimensions. The first dimension represents the solution number, and should have a value of *n* or greater. The second dimension represents the element within the solution *seg*, and should be equal to the segment size.

WriteChild Call

call WriteChild (*selected*, *seg*, *n*, *source*) ;

The WriteChild call assigns values to a selected child solution from within a user crossover operator. The inputs to the WriteChild subroutine are as follows:

- selected* is an array specifying the selected family of solutions. The *selected* array is normally passed into the user subroutine that calls WriteChild, and should be passed unaltered to WriteChild.
- seg* is the segment to which the elements are to be written.
- n* is the child within the family to which the elements are to be written. A value of 1 is for the first child, 2 for the second, and so on.
- source* is an array containing the values to be written.

The WriteChild subroutine is called inside a user crossover operator subroutine to assign to the elements of a selected child solution. It is normally used to complete the action of the crossover operator.

WriteMember Call

call WriteMember (*selected*, *seg*, *source*) ;

The WriteMember call assigns values to a selected solution from within a user objective function or mutation operator. The inputs to the WriteMember subroutine are as follows:

- selected* is an array specifying the selected family of solutions. The *selected* array is normally passed into the user subroutine that calls WriteMember, and should be passed unaltered to WriteMember.
- seg* is the segment to which the elements are to be written.
- source* is an array containing the values to be written.

The WriteMember subroutine is called inside a user objective function or mutation operator subroutine to assign values to the elements of a selected solution. It is normally used to complete the action of the objective function or mutation operator.

Details: GA Procedure

Using Multisegment Encoding

The GA procedure enables you to represent problems with solutions consisting of mixed parameter types by using multisegment encoding. Solutions can contain multiple segments, where each segment is a vector of one particular parameter type. Multiple segments can also be used to store additional information you want to keep for each solution for user objective functions or genetic operators. The utility functions provided by the GA procedure give you full access to read from and write to individual solution segments you define.

Segments are set up with a [SetEncoding call](#). The input parameter to this call is a string consisting of letter-number pairs, with each pair describing the type and number of elements in one segment. The permitted letters and corresponding encodings are as follows:

- R* or *r* specifies real encoding. The elements of the solution segment are real numbers. One common problem where this encoding is used is nonlinear function optimization over the real domain.
- I* or *i* specifies integer encoding. The elements of the solution segment are integers. Examples of where this encoding might be used include assignment problems where the integers represent which resources are assigned to particular tasks, or problems involving real variables that are constrained to be integers.
- B* or *b* specifies Boolean encoding. The elements of the solution consist of binary (0 or 1) bits. This type of encoding might be used, for example, to represent variables in a variable selection problem, or inclusion of particular items in a 0/1 knapsack problem.
- S* or *s* specifies sequence encoding. The segment consists of randomly ordered sequences of integers ranging from 1 to the number of elements. For example, [2, 4, 5, 1, 3] is an example of S5 encoding, as is [5, 3, 2, 1, 4]. Sequence encoding is a natural way to represent routing optimizations like the traveling salesman problem, or any problem optimizing permutations of parameters.

Suppose the problem is to optimize the scheduling of 20 tasks, and for each task you need to choose one machine out of a set of appropriate machines for each task. The natural encoding for that problem could be set up with the following call:

```
call SetEncoding('I20S20');
```

This call specifies a two-segment solution encoding, with segment 1 (I20) an integer vector representing the machine assignment for each task, and segment 2 (S20) representing the sequence of tasks.

When you use multisegment encoding, you must specify the segment parameter in `SetMut` or `SetCross` calls to specify an operator in a segment other than the first one. If you code your own operator subroutines, you can use utility functions provided by the GA procedure to extract and write out values to individual segments of the solution, and routines provided by the GA procedure to perform standard genetic crossover and mutation operations on selected segments. See the section “[Using Standard Genetic Operators and Objective Functions](#)” on page 54 for a discussion of the operators provided by the GA procedure. See the section “[Defining User Genetic Operators](#)” on page 64 and the section “[Defining an Objective Function](#)” on page 68 for details of defining user routines.

Using Standard Genetic Operators and Objective Functions

The GA procedure includes a set of objective functions and crossover and mutation operators, and also enables you to define your own with a subroutine. The standard operators and objectives that are provided for each encoding type are summarized in [Table 3.9](#).

Table 3.9 Standard Genetic Operators for Each Encoding

Encoding	Crossover	Mutation	Objective
real	arithmetic	delta	
	heuristic	null	
	null	uniform	
	simple		
	twopoint		
	uniform		
integer	arithmetic	delta	
	null	null	
	simple	uniform	
	twopoint		
	uniform		
Boolean	simple	null	
	null	uniform	
	twopoint		
	uniform		
sequence	cycle	invert	TSP
	null	null	
	order	swap	
	pmatch		

The following sections describe the standard genetic operators and objectives and how to invoke them.

Crossover Operators

Arithmetic

This operator is defined for real and integer encoding. It treats the solution segment as a vector, and computes offspring of parents P and Q as

$$child1 = aP + (1 - a)Q$$

$$child2 = aQ + (1 - a)P$$

where a is a random number between 0 and 1 generated by the GA procedure. For integer encoding, each component is rounded to the nearest integer. It has the advantage that it always produces feasible offspring for a convex solution space. A disadvantage of this operator is that it tends to produce offspring toward the interior of the search region, so it might not work if the optimum lies on or near the search region boundary. For single-segment encoding, you can specify the use of this operator with the call

```
call SetCross( 'Arithmetic' );
```

For multisegment encoding, you can specify the segment to which the operator should be applied with the call

```
call SetCross( 'Arithmetic', segment );
```

From within a user crossover subroutine, you can use the call

```
call Cross(selected, segment, 'Arithmetic');
```

where *selected* is the selection parameter passed to your subroutine, and *segment* is the segment to which the arithmetic crossover operator is to be applied.

Cycle

This operator is defined for sequence encoding. It produces offspring such that the position of each element value in the offspring comes from one of the parents. For example, consider parents P and Q ,

$$P = [1, 2, 3, 4, 5, 6, 7, 8, 9]$$

$$Q = [8, 7, 9, 3, 4, 1, 2, 5, 6]$$

For the first child, pick the first element from the first parent:

$$child1 = [1, ., ., ., ., ., ., ., .]$$

To maintain the condition that the position of each element value must come from one of the parents, the position of the '8' value must come from P , because the '8' position in Q is already taken by the '1' in $child1$:

$$child1 = [1, ., ., ., ., ., 8, ., .]$$

Now the position of '5' must come from P , and so on until the process returns to the first position:

$$child1 = [1, ., 3, 4, 5, 6, ., 8, 9]$$

At this point, choose the remaining element positions from Q :

$$child1 = [1, 7, 3, 4, 5, 6, 2, 8, 9]$$

For the second child, starting with the first element from the second parent, similar logic produces

$$child2 = [8, 2, 9, 3, 4, 1, 7, 5, 6]$$

This operator is most useful when the absolute position of the elements is of most importance to the objective value. For single-segment encoding, you can specify this operator with the call

```
call SetCross( 'Cycle' );
```

For multisegment encoding, you can specify the segment to which the operator should be applied with the call

```
call SetCross( 'Cycle', segment );
```

From within a user crossover subroutine, you can use the use the call

```
call Cross(selected, segment, 'Cycle');
```

where *selected* is the selection parameter passed to your subroutine, and *segment* is the segment to which the cycle crossover operator is to be applied.

Heuristic

This operator is defined for real encoding. It treats the solution segments as real vectors. It computes the first offspring from two parents *P* and *Q*, where *Q* is the parent with the best objective value, as

$$child1 = a(Q - P) + Q$$

$$child2 = aQ + (1 - a)P$$

where *a* is a random number between 0 and 1 generated by the GA procedure. The first child is a projection, and the second child is a convex combination, as with the arithmetic operator. This operator is unusual in that it uses the objective value. It has the advantage of directing the search in a promising direction, and automatically fine-tuning the search in an area where solutions are clustered. If the solution space has upper and lower bound constraints, the offspring are checked against the bounds, and any component outside its bound is set equal to that bound. The heuristic operator performs best when the objective function is smooth, and might not work well if the objective function or its first derivative is discontinuous.

For single-segment encoding, you can specify this operator with the call

```
call SetCross( 'Heuristic' );
```

For multisegment encoding, you can specify the segment to which the operator should be applied with the call

```
call SetCross( 'Heuristic', segment );
```

From within a user crossover subroutine, you can use the call

```
call Cross(selected, segment, 'Heuristic');
```

where *selected* is the selection parameter passed to your subroutine, and *segment* is the segment to which the heuristic crossover operator is to be applied.

Null

This operator is used to specify that no crossover operator be applied. It is not usually necessary to specify the null operator, except when you want to cancel a previous operator selection. You can specify the null operator with the call

```
call SetCross( 'null' );
```

For multisegment encoding, you can specify a segment to which the operator should be applied with the call

```
call SetCross( 'null', segment );
```

Order

This operator is defined for sequence encoding. It produces offspring by transferring a randomly chosen subsequence of random length and position from one parent, and filling the remaining positions according to the order from the other parent. For parents P and Q , first choose two random cutpoints to define a subsequence:

$$P = [1, 2, |3, 4, 5, 6, |7, 8, 9]$$

$$Q = [8, 7, |9, 3, 4, 1, |2, 5, 6]$$

$$child1 = [., ., 3, 4, 5, 6, ., ., .]$$

$$child2 = [., ., 9, 3, 4, 1, ., ., .]$$

Starting at the second cutpoint and cycling back to the beginning, the elements of Q in order are as follows:

2 5 6 8 7 9 3 4 1

After removing 3, 4, 5 and 6, which have already been placed in $child1$, you have the following entries:

2 8 7 9 1

Placing these back in order starting at the second cutpoint yields the following sequence:

$$child1 = [9, 1, 3, 4, 5, 6, 2, 8, 7]$$

Applying this logic to $child2$ yields the following sequence:

$$child2 = [5, 6, 9, 3, 4, 1, 7, 8, 2]$$

This operator maintains the similarity of the relative order, or adjacency, of the sequence elements of the parents. It is especially effective for circular path-oriented optimizations, such as the traveling salesman problem. For single-segment encoding, you can specify this operator with the call

```
call SetCross( 'Order' );
```

For multisegment encoding, you can specify the segment to which the operator should be applied with the call

```
call SetCross( 'Order', segment );
```

From within a user crossover subroutine, you can use the call

```
call Cross(selected, segment, 'Order' );
```

where *selected* is the selection parameter passed to your subroutine, and *segment* is the segment to which the order crossover operator is to be applied.

Pmatch

The partial match operator is defined for sequence encoding. It produces offspring by transferring a subsequence from one parent, and filling the remaining positions in a way consistent with the position and ordering in the other parent. Start with two parents and randomly chosen cutpoints as indicated:

$$P = [1, 2, |3, 4, 5, 6, |7, 8, 9]$$

$$Q = [8, 7, |9, 3, 4, 1, |2, 5, 6]$$

The first step is to cross the selected subsegments as follows (note that ‘.’ indicates positions yet to be determined):

$$child1 = [., ., 9, 3, 4, 1, ., ., .]$$

$$child2 = [., ., 3, 4, 5, 6, ., ., .]$$

Next, define a mapping according to the two selected subsegments:

$$9-3, \quad 3-4, \quad 4-5, \quad 1-6$$

Then, fill in the positions where there is no conflict from the corresponding parent, as follows:

$$child1 = [., 2, 9, 3, 4, 1, 7, 8, .]$$

$$child2 = [8, 7, 3, 4, 5, 6, 2, ., .]$$

Last, fill in the remaining positions from the subsequence mapping. In this case, for the first child, $1 \rightarrow 6$ and $9 \rightarrow 3$, $3 \rightarrow 4$, $4 \rightarrow 5$, and for the second child, $5 \rightarrow 4$, $4 \rightarrow 3$, $3 \rightarrow 9$, and $6 \rightarrow 1$.

$$child1 = [6, 2, 9, 3, 4, 1, 7, 8, 5]$$

$$child2 = [8, 7, 3, 4, 5, 6, 2, 9, 1]$$

This operator tends to maintain similarity of both the absolute position and relative ordering of the sequence elements, it is useful for a wide range of sequencing problems. For single-segment encoding, you can specify this operator with the call

```
call SetCross( 'Pmatch' );
```

For multisegment encoding, you can specify the segment to which the operator should be applied with the call

```
call SetCross( 'Pmatch', segment );
```

From within a user crossover subroutine, you can use the call

```
call Cross( selected, segment, 'Pmatch' );
```

where *selected* is the selection parameter passed to your subroutine, and *segment* is the segment to which the Pmatch crossover operator is to be applied.

Simple

This operator is defined for integer, real, and Boolean encoding. It has one property, *alpha*. This operator performs the following action: a position k within an encoding of length n is chosen at random, such that $1 \leq k < n$. Then for parents P and Q , the offspring are as follows:

$$child1 = [P_1, P_2, \dots, P_k, Q_{k+1}, Q_{k+2}, \dots, Q_n]$$

$$child2 = [Q_1, Q_2, \dots, Q_k, P_{k+1}, P_{k+2}, \dots, P_n]$$

For integer and real encoding, you can specify an additional *alpha* property of value a , where $0 < a \leq 1$. It modifies the offspring as follows:

$$p_i = aP_i + (1 - a)Q_i, i = k + 1, k + 2, \dots, n$$

$$q_i = aQ_i + (1 - a)P_i, i = k + 1, k + 2, \dots, n$$

$$child1 = [P_1, P_2, \dots, P_k, q_{k+1}, q_{k+2}, \dots, q_n]$$

$$child2 = [Q_1, Q_2, \dots, Q_k, p_{k+1}, p_{k+2}, \dots, p_n]$$

For integer encoding, the elements are then rounded to the nearest integer. For Boolean encoding, the a parameter is ignored, and is effectively 1.

For single-segment encoding, you can specify this operator with the call

```
call SetCross( 'Simple', 'alpha', a );
```

For multisegment encoding, you can specify the segment to which the operator should be applied with the call

```
call SetCross( 'Simple', segment, 'alpha', a );
```

From within a user crossover subroutine, you can use

```
call Cross(selected, segment, 'Simple', a );
```

where *selected* is the selection passed to your subroutine, and *segment* is the segment to which the simple crossover operator is to be applied.

Twopoint

This operator is defined for integer, real, and Boolean encoding of length $n \geq 3$, and has one property, *alpha*. Two positions, $k1$ and $k2$, are chosen at random, such that $1 \leq k1 < k2 < n$. Element values between those positions are swapped between parents. For parents Q and P , the offspring are as follows:

$$child1 = [P_1, P_2, \dots, P_{k1}, Q_{k1+1}, \dots, Q_{k2}, P_{k2+1}, \dots, P_n]$$

$$child2 = [Q_1, Q_2, \dots, Q_{k1}, P_{k1+1}, \dots, P_{k2}, Q_{k2+1}, \dots, Q_n]$$

For integer and real encoding, you can specify an additional *alpha* property of value a , where $0 < a \leq 1$. It modifies the offspring as follows:

$$p_i = aP_i + (1 - a)Q_i, i = k1 + 1, k1 + 2, \dots, k2$$

$$q_i = aQ_i + (1 - a)P_i, i = k1 + 1, k1 + 2, \dots, k2$$

$$child1 = [P_1, P_2, \dots, P_{k1}, q_{k1+1}, \dots, q_{k2}, P_{k2+1}, \dots, P_n]$$

$$child2 = [Q_1, Q_2, \dots, Q_{k1}, p_{k1+1}, \dots, p_{k2}, Q_{k2+1}, \dots, Q_n]$$

Note that small values of a reduce the difference between the offspring and parents. For Boolean encoding, a is always 1. For single-segment encoding, you can specify the use of this operator with the call

```
call SetCross( 'Twopoint', 'alpha', a );
```

For multisegment encoding, you can specify the segment to which the operator should be applied with the call

```
call SetCross( selected, segment, 'Twopoint', 'alpha', a );
```

From within a user crossover subroutine, you can use the call

```
call Cross( selected, segment, 'Twopoint', a );
```

where *selected* is the selection passed to your subroutine, and *seg* is the segment to which the two-point crossover operator is to be applied.

Uniform

This operator is defined for integer, real, and Boolean encoding of length $n \geq 3$, and has two properties, α and p , where $0 < \alpha \leq 1$ and $0 < p \leq 0.5$. For $\alpha = a$ and parents S and T , offspring s and t are generated such that

$$s_i = \begin{cases} aT_i + (1-a)S_i, & \text{with probability } p \\ S_i, & \text{otherwise} \end{cases}$$

$$t_i = \begin{cases} aS_i + (1-a)T_i, & \text{with probability } p \\ T_i, & \text{otherwise} \end{cases}$$

Note that α and p determine how much interchange there is between parents. Lower values of α and p imply less change between offspring and parents. For Boolean encoding, α is always assigned to be 1. If you do not specify α , it defaults to a value of 1. If you do not specify p , it defaults to a value of 0.5.

For single-segment encoding, you can specify the use of this operator with the call

```
call SetCross( 'Uniform', 'alpha', a, 'p', p );
```

For multisegment encoding, you can specify the segment to which the operator should be applied with the call

```
call SetCross( 'Uniform', segment, 'alpha', a, 'p', p );
```

From within a user crossover subroutine, you can use the call

```
call Cross( selected, seg, 'Uniform', a, p );
```

where *selected* is the selection parameter passed to your crossover subroutine, and *seg* is the segment to which the uniform crossover operator is to be applied. In both cases, if the encoding is Boolean, the a parameter is ignored, and a is effectively 1.

Mutation Operators

Delta

This operator is defined for integer and real encoding. It has two properties, *delta* and *nchange*. It first chooses *n* elements of the solution at random, where *n* is the value of the *nchange* property, and then perturbs each chosen element by a fixed amount, set by *d*, the value of the *delta* property. *d* must be an array with the same length as the encoding. A randomly chosen element *k* of the solution *S* is modified such that

$$S_k \in \{S_k - d_k, S_k + d_k\}$$

If upper and lower bounds are specified with a [SetBounds](#) call, then S_k is adjusted as necessary to fit within the bounds. This operator gives you the ability to fine-tune the search by modifying the magnitude of the *delta* property. One possible strategy is to start with larger *d* values, and then reduce them as the search progresses and begins to converge to an optimum. This operator is also useful if the optimum is known to be on or near a boundary, in which case *d* can be set large enough to always perturb the solution element to a boundary. For single-segment encoding, you can specify this operator with the call

```
call SetMut ( 'delta', 'nchange', n, 'delta', d ) ;
```

For multisegment encoding, you can specify the segment to which the operator should be applied with the call

```
call SetMut ( 'delta', segment, 'nchange', n, 'delta', d ) ;
```

You can invoke this operator on a solution segment from within a user mutation subroutine with the call

```
call Mutate ( selected, segment, 'delta', d, n ) ;
```

where *selected* is the selection parameter passed into the subroutine.

Invert

This operator is defined for sequence encoding. It picks two locations at random and reverses the order of elements between them. This operator is most often applied to the traveling salesman problem. For single-segment encoding, you can specify this operator with the call

```
call SetMut ( 'invert' ) ;
```

For multisegment encoding, you can specify the segment to which the operator should be applied with the call

```
call SetMut ( 'invert', segment ) ;
```

You can invoke this operator on a solution segment from within a user mutation subroutine with the call

```
call Mutate ( selected, segment, 'invert' ) ;
```

where *selected* is the selection parameter passed into the subroutine.

Null

This operator is used to specify that no mutation operator be applied. It is not usually necessary to specify the null operator, except when you want to cancel a previous operator selection. You can specify the null operator with the call

```
call SetMut ( 'null' ) ;
```

For multisegment encoding, you can specify a segment to which the operator should be applied with

```
call SetMut ( 'null', segment ) ;
```

Swap

This operator is defined for sequence problem encoding. It picks two random locations in the solution vector and swaps their values. You can also specify that multiple swaps be made for each mutation, by setting the *nswap* property. For single-segment encoding, you can specify the swap operator with the call

```
call SetMut ( 'swap', 'nswap', n ) ;
```

where *n* is the number of swaps for each mutation. For multisegment encoding, you can specify the segment to which the operator should be applied with the call

```
call SetMut ( 'swap', segment, 'nswap', n ) ;
```

You can invoke this operator on a solution segment from within a user mutation subroutine with the call

```
call Mutate ( selected, segment, 'swap', n ) ;
```

where *selected* is the selection parameter passed into the subroutine.

Uniform

This operator is defined for Boolean encoding or for real or integer encoding with upper and lower bounds specified with a [SetBounds](#) call. To apply this operator, a position *k* is randomly chosen within the solution *S*, and *S_k* is modified to a random value between the upper and lower bounds for element *k*. There are two properties you can specify, *nchange* and *pchange*. If you set the *nchange* property to *n*, the operator is applied to *n* locations. If you set the *pchange* property to a value *p*, where $0 < p < 1$, the operator is applied at each position with probability *p*. Only the last property setting is active, overriding any previous settings. If no property is set, a default *nchange* value of 1 is used. You can specify this operator with one of the following two calls:

```
call SetMut ( 'uniform', 'nchange', n ) ;
```

```
call SetMut ( 'uniform', 'pchange', p ) ;
```

For multisegment encoding, you can specify the segment to which the operator should be applied with the call

```
call SetMut ( 'uniform', segment, 'nchange', n ) ;
```

You can invoke this operator on a solution segment from within a user mutation subroutine by using one of the following statements, where *selected* is the selection parameter passed into the subroutine:

```
call Mutate ( selected, segment, 'uniform', n ) ;
```

```
call Mutate ( selected, segment, 'uniform', p ) ;
```

This operator can prove especially useful in early stages of the optimization, since it tends to distribute solutions widely across the search space, and to avoid premature convergence to a local optimum. However, in later stages of an optimization, when the search needs to be fine-tuned to home in on an optimum, the uniform operator can hinder the optimization.

Objective Functions

TSP

The TSP objective calculates the value of the traveling salesman problem objective. It has one property, *distances*, which is a two-dimensional array representing distances between locations. If $distances = d$, then $d[i, j]$ is the distance between location i and location j .

Defining a User Fitness Comparison Routine

Most of the selection techniques used in this procedure require comparisons of fitness to decide which solutions should be propagated into the next generation. Normally fitness is directly related to the solution objective value, with the better objective value indicating the higher fitness. However, sometimes it is useful to base fitness on multiple criteria, not just a single calculated objective value. Doing so enables you to solve multiobjective optimization problems, or to find optima that also satisfy secondary objectives. For example, for a multiobjective problem you might want to find a set of construction schedules that simultaneously minimize completion time and minimize cost, so that you can examine the tradeoff between the two objectives and choose a schedule that best meets your needs. You might also have a scheduling problem where there is more than one solution that meets your primary objective, but you would like to converge on the one that also provides adequate worker breaks, or minimizes overtime. One single-valued objective function might not be adequate to express these preferences.

You can define a function and designate that function to be used for comparing the fitness of two competing solutions with the call

```
call SetCompareRoutine( 'routine' );
```

where *routine* is the name of the function you have defined. The function name must be a quoted string. The first parameter of the function, designated the *selection* parameter, must be a numeric array. The procedure calls your compare function whenever it needs to compare the fitness of two solutions. Your function can also have an arbitrary number of additional parameters, whose names and data types should match variables defined in the global portion of your input. When the procedure calls your function, the *selection* parameter will be filled in with information identifying which solutions are to be compared, and any other parameters will be filled in with the corresponding global symbol values. Your function should not alter the selection parameter in any way, but pass it unchanged into a [ReadCompare Call](#) to obtain the solution elements your program needs to compare the fitness. If solution 1 has higher fitness than solution 2, then your function should return a positive number. If solution 2 is more fit than solution 1, then a negative number should be returned. If the two solutions have the same fitness then a value of 0 might be returned. Following is a simple example of a fitness comparison function.

```
/* This function is designed to maximize a
 * primary and secondary objective. The objectives
 * are stored in segment 2 of the encoding. If
 * the difference in primary objective is less than
 * the value of delta, then the secondary objective
 * is used to determine the fitness
 */
function compare2(selected[*], delta);

/* arrays to hold solution elements */
```

```

array member1[2] /nosym;
array member2[2] /nosym;

/* read segment 2 of solution 1 into member1 */
call ReadCompare(selected, 2, 1, member1);

/* read segment 2 of solution 2 into member2 */
call ReadCompare(selected, 2, 2, member2);

/* element 1 contains the primary objective value, */
/* element 2 contains a secondary objective value */

/* if objective 1 is nearly the same, then use */
/* objective 2 to compare the fitness */
if( abs(member1[1] - member2[1]) < delta) then do;
    /* primary objectives are nearly the same, check secondary */
    if(member1[2] > member2[2]) then
        return(1);
    if(member2[2] > member1[2]) then
        return(-1);
end;

/* base fitness on primary objective */
if(member1[1] > member2[1]) then
    return(1);
if(member2[1] > member1[1]) then
    return(-1);

/* all objectives are the same, return 0 */
return(0);
endsub;

/* in global scope of the input */
delta = 0.01;
call SetCompareRoutine('compare2');

```

For an example of a comparison routine used in a multiobjective optimization, see [Example 3.3](#).

Defining User Genetic Operators

You can define new genetic operators with subroutines. The GA procedure calls your subroutine when it is necessary to perform mutation or crossover operations. You can designate that a subroutine be used for crossover with the call

```
call SetCrossRoutine('routine');
```

where *routine* is the name of the subroutine you have defined. Similarly, you can designate a subroutine for the mutation operator with the call

```
call SetMutRoutine('routine');
```

The subroutine name must be a quoted string. The first parameter of the crossover or mutation subroutine you define must be a numeric array. When the GA procedure calls your subroutine, it passes information

in the first parameter, referred to as the *selection* parameter, which designates the selected members for the operation. You should not alter the selection parameter in any way, but pass it unchanged into special utility routines provided by the GA procedure in order to obtain the solution elements and write them to the selected members. You can define as many other parameters to your subroutine as you need; they are filled in with values from variables of the same name created in the global portion of your program. Any array parameters must be numeric and of the type /NOSYMBOLS.

For a crossover subroutine, use the [ReadParent](#) call to get the elements of the selected parents into arrays that you can then manipulate with programming statements. The results can be written to the designated offspring with a [WriteChild](#) call. The following code is an example of a crossover subroutine. The subroutine creates two new offspring from two selected parents by switching the odd-numbered elements between the two parents.

```
/* single-segment integer encoding of size 10 */
call SetEncoding('I10');

/* encoding size is 10 */
n = 10;

subroutine swapodd(selected[*], n);
  array child1[1] /nosym;
  array child2[1] /nosym;

  /* reallocate child arrays to right size */
  call dynamic_array(child1,n);
  call dynamic_array(child2,n);

  /* read segment 1 from parent 1 into child1 */
  call ReadParent(selected, 1, 1, child1);

  /* read segment 1 from parent 2 into child2 */
  call ReadParent(selected, 1, 2, child2);

  /* swap the odd elements in the solution */
  do i = 1 to n by 2;
    temp = child1[i];
    child1[i] = child2[i];
    child2[i] = temp;
  end;

  /* write offspring out to selected children */
  call WriteChild(selected, 1, 1, child1);
  call WriteChild(selected, 1, 2, child2);
endsub;

/* designate swapodd as the crossover routine */
call SetCrossRoutine('swapodd');
```

The next sample program illustrates a crossover routine that might be used for multisegment mixed integer and sequence encoding. The subroutine uses the standard Simple crossover operator for the integer segment, and the Pmatch operator for the sequence-encoded segment.

```

/* Solution has 2 segments, integer I5 and sequence S5 */
call SetEncoding('I5S5');

/* alpha parameter for Simple crossover operator */
alpha = 1;

subroutine mixedIS(selected[*], alpha);

    /* execute simple operator on segment 1 */
    call Cross(selected, 1, 'Simple', alpha);

    /* execute pmatch operator on segment 2 */
    call Cross(selected, 2, 'Pmatch');

endsub;

call SetCrossRoutine('mixedIS');

```

For a mutation subroutine, use a [ReadMember](#) call to obtain the elements of the solution selected to be mutated, and use a [WriteMember](#) call to write the mutated elements back to the solution. For example, the following statements define a mutation subroutine that swaps two adjacent elements at a randomly chosen position in a sequence:

```

/* Solution has 1 segment, sequence S10 */
call SetEncoding('S10');

n = 10;

subroutine swap2(selected[*], n);

    /* declare an array for working memory */
    array member[1] /nosym;

    /* allocate array to required length */
    call dynamic_array(member, n);

    /* read segment 1 of selected member into array */
    call ReadMember(selected,1,member);

    /* generate random number between 0 and 1 */
    r = rand('uniform');

    /* convert r to integer between 1 and n-1 */
    i = int(r * (n - 1)) + 1;

    /* swap element values */
    temp = member[i];
    member[i] = member[i+1];
    member[i+1] = temp;

    /* write result back out to solution */
    call WriteMember(selected,1,member);

```

```
endsub;

/* Set the mutation routine to swap2 */
call SetMutRoutine('swap2');
```

Defining a User Update Routine

The GA procedure enables you to define a routine that is to be called at each generation in the optimization process, after the designated objective function has been evaluated for each solution, before the selection process takes place. Some of the tasks that can be handled with an update routine include evaluating global solution fitness criteria, logging intermediate optimization progress, evaluating termination criteria and stopping the optimization, modifying the properties of the genetic operators to intensify or diversify the search process, or even to reinitialize portions of the solution population. You can designate a user update function with the call

```
call SetUpdateRoutine('name');
```

where *name* is the name of the routine you have defined.

The parameters defined for the update routine must correspond to variables of the same name established in the global portion of your procedure input. If you desire to update a parameter and have that update retained across generations, that parameter must also be declared in an OUTARGS statement in the update routine definition. The following program snippet illustrates how a user might specify an update routine that monitors the best objective value at each generation and terminates the optimization when no improvement is seen after a specified number of iterations.

```
subroutine no_improvement_terminator(iteration,
                                     saved_value,
                                     nsame,
                                     same_limit,
                                     populationSize);
outargs iteration, saved_value, nsame;

array objValues[1] /nosym;

/* dynamically allocate array to fit populationSize */
call dynamic_array(objValues, populationSize);

/* read in current objective values */
call GetObjValues(objValues, populationSize);

/* find best value */
current_best = objValues[1];
do i = 2 to populationSize;
  /* for a minimization problem, use < here */
  if(objValues[i] > current_best) then
    current_best = objValues[i];
end;
```

```

if iteration > 0 then do;
  /* for a minimization problem, use < here */
  if current_best > saved_value then do;
    /* reset same value counter */
    nsame = 1;
    saved_value = current_best;
  end;
  else do;
    /* increment same value counter */
    nsame = nsame + 1;
    /* if counter equals limit, then make this the last generation */
    if nsame >= same_limit then
      call ContinueFor(0);
    end;
  end;
end;
else do;
  saved_value = current_best;
  nsame = 1;
end;

iteration = iteration + 1;

endsub;

iteration = 0;
saved_value = 0;
nsame = 0;
/* terminate when no improvement after 30 generations */
same_limit = 30;
populationSize = 100;
call SetUpdateRoutine('no_improvement_terminator');

```

From within a user update routine you can use a [GetObjValues](#) call to get the current solution objective values, a [GetSolutions](#) call to get the current solution population, an [UpdateSolutions](#) call to reset solution values, a [ReEvaluate](#) call to recompute objective values, or an [Initialize](#) call to reinitialize and resize the solution population.

Defining an Objective Function

The GA procedure enables you to specify your objective to be optimized with a function you create, or as a standard objective function that the GA procedure provides. Currently the only standard objective you can specify without writing an objective function is the traveling salesman problem, which can be specified with a [SetObj](#) call. In the future, other objective functions will be added. You can designate a user objective function with the call

```
call SetObjFunc('name', minmax);
```

where *name* is the name of the function you have defined, and *minmax* is set to 0 to specify a minimum or 1 to specify a maximum.

A user objective function must have a numeric array as its first parameter. When the GA procedure calls your function, it passes an array in the first parameter that specifies the selected solution, which is referred to

as the *selection* parameter. The selection parameter must not be altered in any way by your function. Your function should pass the selection parameter to a [ReadMember call](#) to read the elements of the selected solution into an array. Your function can then access this array to compute an objective value, which it must return. As with the genetic operator routines, you can define additional arguments to your objective function, and the GA procedure passes in variables with corresponding names that you have created in your global program. For example, the following statements set up an objective function that minimizes the sum of the squares of the solution elements:

```
call SetEncoding('R5');

n = 5;

function sumsq(selected[*], n);

    /* set up a scratch array to hold solution elements */
    array x[1] /nosym;

    /* allocate x to hold all solution elements */
    call dynamic_array(x, n);

    /* read members of the selected solution into x */
    call ReadMember(selected, 1, x);

    /* compute the sum of the squares */
    sum = 0;
    do i = 1 to n;
        sq = x[i] * x[i];
        sum = sum + sq;
    end;

    /* return the objective value */
    return(sum);
endsub;

call SetObjFunc('sumsq', 0);
```

In this example, the function `SUMSQ` is defined, and the [SetObjFunc call](#) establishes it as the objective function. The 0 for the second parameter of the `SetObjFunc` call indicates that the objective should be minimized. Note that the second parameter to the `SUMSQ` function, *n*, is defined in the procedure, and the value assigned to it there is passed into the function.

Defining a User Initialization Routine

For problems with simple constant bounds or simple sequencing problems it is not necessary to define a user initialization subroutine; simply specify 'DEFAULT' in the [Initialize call](#). Defining a routine is necessary only if you need to satisfy more complicated constraints or apply some initial heuristics or local optimizations. A user initialization routine is specified with an `Initialize` call, as follows:

```
call Initialize('name', size);
```

where *name* is the name of your initialize routine. The first parameter of the subroutine you define must be a numeric array. When the GA procedure calls your subroutine, it passes information in the first parameter, referred to as the *selection* parameter, which designates the member selected for initialization. Your subroutine should generate one solution and write the values of the solution elements with a [WriteMember](#) call, using the selection parameter passed to your subroutine. The random number functions from BASE SAS are available to your subroutine, if needed. You can define as many other parameters to your subroutine as you need; they are filled in with values from variables of the same name created in your global program. The array used to write the generated solution to the population must be numeric and declared with the /NOSYMBOLS option, as well as any arrays passed as parameters into your subroutine.

The following sample statements illustrate how to define an initialization routine. The feasible region is a triangle with vertices (0,0), (0,1) and (1,1).

```
call SetEncoding('R2');

/* set vertices of triangle (0,0), (0,1), and (1,1) */
array vertex1[2] /nosym (0,0);
array vertex2[2] /nosym (0,1);
array vertex3[2] /nosym (1,1);

subroutine triangle(selected[*], vertex1[2], vertex2[2], vertex3[2]);

array x[2] /nosym;

/* select 3 random numbers 0 < r < 1 */
r1 = rand('uniform');
r2 = rand('uniform');
r3 = rand('uniform');

/* normalize so r1 + r2 + r3 = 1 */
sumr = r1 + r2 + r3;
r1 = r1 / sumr;
r2 = r2 / sumr;
r3 = r3 / sumr;

/* form a convex combination of vertices in x */
do i = 1 to 2;
    x[i] = r1 * vertex1[i] + r2 * vertex2[i] + r3 * vertex3[i];
end;

/* write x out to the selected population member, to segment 1 */
call WriteMember(selected, 1, x);
endsub;

[other programming statements]

call Initialize('triangle',100);
```

In this example, the triangle initialization subroutine generates a solution that is a random convex combination of three points, which places it in the interior of the triangular region defined by the points. Note the use of the BASE SAS `RAND()` function to get random numbers uniformly distributed between 0 and 1. The random numbers are then normalized so that their sum is 1. In the loop, they are used to compute

a convex linear combination of the vertices, and the `WriteMember` call writes the solution to the selected population member. The encoding specified a single segment, so the `WriteMember` call specifies segment 1 as the target. When the GA procedure executes the `Initialize` call, it executes the triangle routine 100 times, once for each member of the initial population.

Specifying the Selection Strategy

There are a number of different strategies that can be employed to select the most fit solutions from a population to be propagated into the next solution generation. Regardless of the strategy chosen, the user must try to set the selection properties to maintain a productive balance between selection pressure and preservation of diversity. Enough selective pressure must be maintained to drive the algorithm toward an optimum in a reasonable time, but too much selective pressure will eliminate the diversity of the population too quickly and lead to premature convergence to a suboptimal solution. One effective technique is to start the optimization process at a lower selective pressure, and increase it as the optimization continues. You can easily do this in the GA procedure by modifying the selection strategy with a `SetProperty` call inside an `update` routine. A description of the selection strategies and their properties follows.

Tournament Selector

This selector has one property, *size*. The selector repeats the following process until the next generation of solutions has been specified: randomly choose a small subset of the current population, and then select the most fit from that subset. The selection pressure is controlled by the size of the subset chosen, specified by the *size* property. Tournament sizes from 2 to 10 have been successfully applied to various genetic algorithm optimizations, with sizes over 4 or 5 considered to represent strong selective pressure. If the *size* property is not set by the user, it defaults to a value of 2. The `Duel` selector is a modification of this selector that permits further lowering of selection pressure.

With this selector the objective value is normally used to compare the fitness of competing solutions, with better objective values corresponding to higher fitness. However, there are techniques for multiple objective optimization and constraint handling that use more complicated criteria than just the objective value. See the section “[Optimizing Multiple Objectives](#)” on page 73 for further discussion of this topic. You can set your own fitness comparison routine to implement those techniques with the `SetCompareRoutine` call, and that routine will be used by the tournament selector to compare solution fitness. For a simple single-objective problem you normally do not need to specify a special fitness comparison routine, unless you have a secondary objective you might want to use to break a tie.

Duel Selector

This selector has one property, *bprob*. It operates in the same manner as the `Tournament` selector with tournament size 2, except it permits you to reduce selection pressure further by specifying a probability, *bprob*, for selecting the most fit solution from the pair. The *bprob* property is assigned a default value of 0.8, but you can change it to a value between 0.5 (corresponding to pure random selection) and 1. By default, solution fitness is compared by comparing the solution objective values. However, you can also set your own fitness comparison routine with the `SetCompareRoutine` call. See the `tournament` selector for a discussion of when you would need to specify a special fitness comparison routine.

Incorporating Heuristics and Local Optimizations

It is often effective to combine the genetic algorithm technique and other local optimizations or heuristic improvements. This can be done within the GA procedure by incorporating a local optimization into a user objective function and returning an improved objective value. Either your user objective function can replace the original solution with the optimized one, or you can leave the solution unchanged, replacing it with the optimized one only at the final iteration.

Replacing the original solution with the locally optimized one speeds convergence, but it also increases the risk of converging prematurely. If you choose to do so, you can modify the solution by writing the changed solution back to the population with a [WriteMember call](#). You could also consider replacing the original solution with some probability p . For some problems, values of p from 0.05 to 0.15 have been shown to significantly improve convergence while avoiding premature convergence to a local optimum. This technique is illustrated in [Example 3.1](#).

Handling Constraints

Practical optimization problems usually involve constraints, which can make the problems harder to solve. Constraints are handled in genetic algorithms in several ways.

Encoding Strategy

The simplest approach is to set the problem encoding, genetic operators, and initialization such that the constraints are automatically satisfied. Fixed constant bounds are easily handled in this manner in the GA procedure with the [SetBounds call](#). The default initialization process and genetic operators provided by the GA procedure automatically respect bounds specified in this manner. For some types of constraints, you might be able to create a direct mapping from a constrained solution domain to a second domain with simple constant bounds. You could then define your genetic operator to map the solution into the second domain, apply one of the standard genetic operators, and then map the result back to the original domain.

If the problem contains equality constraints, you should try to transform the problem to eliminate the equality constraints and reduce the number of variables. This strategy is opposite from what is usually done in linear programming, where inequality constraints are turned into equality constraints by the addition of slack variables.

Repair Strategy

If the constraints are more complex and cannot be easily satisfied automatically by the genetic operators, you might be able to employ a repair strategy: check the solution and modify it to satisfy the constraints. The check and repair can be done in a user genetic operator when the solution is generated, or it can be done in the evaluation phase in a user objective function. Possible strategies for making a repair inside an objective function include projecting the solution onto the constraint boundary; while inside a genetic operator you might try adjusting an operator parameter until the constraint is satisfied. If you do the repair in the objective function, you should compute the objective value after performing the repair. You can write the repaired solution back out to the population with a [WriteMember call](#) from a user objective function or mutation subroutine, and with a [WriteChild call](#) from within a crossover subroutine. [Example 3.2](#) illustrates the use of the repair strategy.

Penalty Strategy

Another technique is to permit solutions to violate constraints, but also to impose a fitness penalty that causes the population to evolve toward satisfying constraints as well as optimizing the objective. One way of employing this strategy is to simply add a penalty term to the objective function, but this approach should be used with care, because it is not always obvious how to construct the penalty function in a way that does not bias the optimization of the desired objective.

Direct Comparison Strategy

Using tournament selection opens another possibility for handling constraints. Define a fitness comparison routine (designated in a [SetCompareRoutine call](#)) that employs the following logic:

- 1 If neither solution is feasible, choose the one closest to satisfying the constraints.
- 2 If one solution is feasible, and the other is not, choose the feasible one.
- 3 If both solutions are feasible, choose the one with the best objective value.

This strategy has the advantage that the objective function does not have to be calculated for infeasible solutions. To implement this method, you need to provide a measure of constraint violation and compute it in a user objective function; this value can be used in the first comparison step outlined previously. For linear constraints, the GA procedure provides the [EvaluateLC call](#) for this purpose. The technique works best when the solution space normally contains a significant number of solutions that satisfy the constraints. Otherwise it is possible that a single feasible solution might quickly dominate the population. In such cases, a better approach might be the following Bicriteria Comparison Strategy.

Bicriteria Comparison Strategy

A variation of the direct comparison strategy that has proved effective in many applications is the multiobjective, bicriteria approach. This strategy involves adding a second objective function, which is the magnitude of the constraint violation. Based on the original and constraint violation objective functions, a Pareto-optimal set of solutions is evolved in the population, and the Pareto-optimal set is evolved toward zero constraint violation. This technique is illustrated in [Example 3.3](#). See the section “[Optimizing Multiple Objectives](#)” on page 73 for a full discussion of Pareto optimality and how to apply this technique.

Optimizing Multiple Objectives

Many practical optimization problems involve more than one objective criteria, where the decision maker needs to examine trade-offs between conflicting objectives. With traditional optimization methods, these problems are often handled by aggregating multiple objectives into a single scalar objective, usually accomplished by some linear weighting of the multiple criteria. Other approaches involve turning objectives into constraints. One disadvantage of this strategy is that many separate optimizations with different weighting factors or constraints need to be performed to examine the trade-offs between different objectives. Genetic algorithms enable you to attack multiobjective problems directly, in order to evolve a set of solutions in one run of the optimization process instead of solving multiple separate problems.

This approach seeks to evolve the Pareto-optimal set: the set of solutions such that for each solution, all the objective criteria cannot be simultaneously improved. This is expressed mathematically by the concept of Pareto optimality. A Pareto-optimal set is the set of all nondominated solutions, according to the following definition of *dominated*:

For an n -objective minimizing optimization problem, for each objective function f_i , a solution p is *dominated* by q if

$$f_i(p) \geq f_i(q) \text{ for all } i = 1, \dots, n \text{ and } f_j(p) > f_j(q) \text{ for some } j = 1, \dots, n$$

The following is one strategy that can be employed in the GA procedure to evolve a set of Pareto-optimal solutions to a multiobjective optimization problem:

user objective function: Define and specify a user objective function in a [SetObjFunc call](#) that computes each of the objective criteria and stores the objective values in one single solution segment.

user update routine: Define and specify a user update routine in a [SetUpdateRoutine call](#) that examines the entire solution population and marks those in the Pareto-optimal set. This can be done with the [MarkPareto call](#) provided by the GA procedure. Also, set the *elite* parameter equal to the number of Pareto-optimal solutions found.

selection criteria: Define a fitness comparison routine that favors the least dominated solutions, and designate it in a [SetCompareRoutine call](#). For selecting between two solutions when neither solution dominates the other, your routine can check a secondary criterion to direct the search to the area of ultimate interest.

The multiple objective values are recorded in one segment to enable the use of the [MarkPareto call](#) provided by the GA procedure. Setting the *elite* selection parameter to the size of the Pareto-optimal set, in conjunction with the comparison criteria, guarantees that the Pareto-optimal set in each generation is preserved to the next.

The secondary comparison criterion can be used to ensure that the final Pareto-optimal set is distributed in the area of ultimate interest. For example, for the Bicriteria Constraint Strategy described previously, the actual area of interest is where there is zero constraint violation, which is the second objective. The secondary comparison criterion in that case is to minimize the value of the constraint violation objective. After enough iterations, the population should evolve to the point that the best solution to the bicriteria problem is also the best solution to the original constrained problem, and the other Pareto-optimal solutions can be examined to analyze the sensitivity of the optimum solution to the constraints. For other types of problems, you might need to implement a more complicated secondary comparison criterion to avoid “crowding” of solutions about some arbitrary point, and ensure the evolved Pareto-optimal set is distributed over a range of objective values.

Examples: GA Procedure

Example 3.1: Traveling Salesman Problem with Local Optimization

This example illustrates the use of the GA procedure to solve a traveling salesman problem (TSP); it combines a genetic algorithm with a local optimization strategy. The procedure finds the shortest tour of 20

locations randomly oriented on a two-dimensional x - y plane, where $0 \leq x \leq 1$ and $0 \leq y \leq 1$. The location coordinates are input in the following DATA step:

```
/* 20 random locations for a Traveling Salesman Problem */
data locations;
    input x y;
    datalines;
0.0333692 0.9925079
0.6020896 0.0168807
0.1532083 0.7020444
0.3181124 0.1469288
0.1878440 0.8679120
0.9786112 0.4925364
0.7918010 0.7943144
0.5145329 0.0363478
0.5500754 0.8324617
0.3893757 0.6635483
0.9641841 0.6400201
0.7718126 0.5463923
0.7549037 0.4584584
0.2837881 0.7733415
0.3308411 0.1974851
0.7977221 0.1193149
0.3221207 0.7930478
0.9201035 0.1186234
0.2397964 0.1448552
0.3967470 0.6716172
;
```

First, the GA procedure is run with no local optimizations applied:

```
proc ga data1 = locations seed = 5554;
    call SetEncoding('S20');
    ncities = 20;
    array distances[20,20] /nosym;
    do i = 1 to 20;
        do j = 1 to i;
            distances[i,j] = sqrt((x[i] - x[j])**2 + (y[i] - y[j])**2);
            distances[j,i] = distances[i,j];
        end;
    end;
    call SetObj('TSP',0,'distances', distances);
    call SetCross('Order');
    call SetMut('Invert');
    call SetMutProb(0.05);
    call SetCrossProb(0.8);
    call SetElite(1);
    call Initialize('DEFAULT',200);
    call ContinueFor(140);
run;
```

The PROC GA statement uses a DATA1= option to get the contents of the locations data set, which creates array variables x and y from the corresponding fields of the data set. A solution will be represented as a circular tour, modeled as a 20-element sequence of locations, which is set up with the [SetEncoding call](#). The

array variable *distances* is created, and the loop initializes *distances* from the *x* and *y* location coordinates such that *distances*[*i*, *j*] is the Euclidean distance between location *i* and *j*. Next, the **SetObj** call specifies that the GA procedure use the included TSP objective function with the *distances* array. Then, the genetic operators are specified, with **SetCross** for the crossover operator and **SetMut** for the mutation operator. Since the crossover probability and mutation probability are not explicitly set in this example, the default values of 1 and 0.05, respectively, are used. The selection parameters are not explicitly set (with **SetSel** and **SetElite** calls), so by default, a tournament of size 2 is used, and an *elite* parameter of 1 is used. Next, the **Initialize** call specifies default initialization (random sequences) and a population size of 100. The **ContinueFor** call specifies a run of 220 iterations. This value is a result of experimentation, after it was determined that the solution did not improve with more iterations. The output of this run of PROC GA is given in **Output 3.1.1**.

Output 3.1.1 Simple Traveling Salesman Problem

PROC GA Optimum Values		
Objective		
3.7465311323		
Solution		
Element	Value	
1	12	
2	13	
3	18	
4	16	
5	2	
6	8	
7	15	
8	4	
9	19	
10	3	
11	1	
12	5	
13	14	
14	17	
15	10	
16	20	
17	9	
18	7	
19	11	
20	6	

The following program illustrates how the problem can be solved in fewer iterations by employing a local optimization. Inside the user objective function, before computing the objective value, every adjacent pair of cities in the tour is checked to determine if reversing the pair order would improve the objective value. For a pair of locations S_i and S_{i+1} , this means comparing the distance traversed by the subsequence $\{S_{i-1}, S_i, S_{i+1}, S_{i+2}\}$ to the distance traversed by the subsequence $\{S_{i-1}, S_{i+1}, S_i, S_{i+2}\}$, with appropriate wrap-around at the endpoints of the sequence. If the distance for the swapped pair is smaller than the original pair, then the reversal is done, and the improved solution is written back to the population.

```

proc ga data1 = locations seed = 5554;
  call SetEncoding('S20');
  ncities = 20;
  array distances[20,20] /nosym;
  do i = 1 to 20;
    do j = 1 to i;
      distances[i,j] = sqrt((x[i] - x[j])**2 + (y[i] - y[j])**2);
      distances[j,i] = distances[i,j];
    end;
  end;
  /* Objective function with local optimization */
  function TSPSwap(selected[,ncities,distances[,*]]);
  array s[1] /nosym;
  call dynamic_array(s,ncities);
  call ReadMember(selected,1,s);
  /* First try to improve solution by swapping adjacent cities */
  do i = 1 to ncities;
    city1 = s[i];
    inext = 1 + mod(i,ncities);
    city2 = s[inext];
    if i=1 then
      before = s[ncities];
    else
      before = s[i-1];
      after = s[1 + mod(inext,ncities)];
    if (distances[before,city1]+distances[city2,after]) >
      (distances[before,city2]+distances[city1,after]) then do;
      s[i] = city2;
      s[inext] = city1;
    end;
  end;
  end;
  call WriteMember(selected,1,s);
  /* Now compute distance of tour */
  distance = distances[s[ncities],s[1]];
  do i = 1 to (ncities - 1);
    distance + distances[s[i],s[i+1]];
  end;
  return(distance);
endsub;
call SetObjFunc('TSPSwap',0);
call SetCross('Order');
call SetMut('Invert');
call SetMutProb(0.05);
call SetCrossProb(0.8);
call SetElite(1);
call Initialize('DEFAULT',200);
call ContinueFor(35);
run;

```

The output after 85 iterations is given in [Output 3.1.2](#).

Output 3.1.2 Traveling Salesman Problem with Local Optimization

PROC GA Optimum Values		
Objective		
3.7465311323		
Solution		
Element	Value	
1	13	
2	18	
3	16	
4	2	
5	8	
6	15	
7	4	
8	19	
9	3	
10	1	
11	5	
12	14	
13	17	
14	10	
15	20	
16	9	
17	7	
18	11	
19	6	
20	12	

Since all tours are circular, the actual starting point does not matter, and this solution is equivalent to that reached with the simple approach without local optimization. It is reached after only 85 iterations, versus 220 with the simple approach.

NOTE: See the “Examples” section in Chapter 2, “The OPTNET Procedure” (*SAS/OR User’s Guide: Network Optimization Algorithms*), for an example of how to use PROC OPTNET to solve the TSP.

Example 3.2: Nonlinear Objective with Constraints Using Repair Mechanism

This example illustrates the use of a repair mechanism to satisfy problem constraints. The problem is to minimize the six-hump camel-back function (Michalewicz 1996, Appendix B):

$$f(x) = \left(4 - 2.1x_1^2 + \frac{x_1^4}{3}\right)x_1^2 + x_1x_2 + (-4 + 4x_2^2)x_2^2$$

where x is within the triangle with vertices $V_1 = (-2, 0)$, $V_2 = (0, 2)$, and $V_3 = (2, -2)$. The problem formulation takes advantage of the fact that all feasible solutions can be expressed as a convex combination of V_1 , V_2 , and V_3 in the following form:

$$x = aV_1 + bV_2 + cV_3$$

In this form, a , b , and c are nonnegative coefficients and satisfy the following linear equality constraint:

$$a + b + c = 1$$

Therefore, the solution encoding 'R3' is used with the three elements corresponding to the values of a , b , and c . Note that this strategy can be generalized to any solution domain that can be specified by a convex hull. An additional 'R2' segment is also created to store the corresponding x value. In the following program, **FUNCTION SIXHUMP** computes the objective value. Lower and upper bounds of 0 and 1, respectively, are used to ensure that solution elements are nonnegative. Violations of the linear equality constraint are fixed by a simple repair strategy, implemented in **FUNCTION SIXHUMP**: the sum of the solution elements is computed, and each element is divided by the sum, so that the sum of the new elements is 1. For the special case of all elements equal to 0, equal values are assigned to the solution elements.

```

proc ga seed = 555;
call SetEncoding('R3R2');
npoints = 3;
array cvxhull[3,2] /nosym ( -2 0
                           0 2
                           2 -2 );

/* Objective function */
function sixhump(selected[,],cvxhull[,],npoints);
  /* Function has global minimum value of -1.0316
   * at x = {-0.0898  0.7126} and
   *      x = { 0.0898 -0.7126}
   */
  array w[1] /nosym;
  call dynamic_array(w,npoints);
  array x[2] /nosym;
  call ReadMember(selected,1,w);
  /* make sure that weights add up to 1 */
  sum = 0;
  do i = 1 to npoints;
    sum + w[i];
  end;
  /* if all weights 0, then reinitialize */
  if sum=0 then do;
    sum = npoints;
    do i = 1 to npoints;
      w[i] = 1;
    end;
  end;
  /* re-normalize weights */
  do i = 1 to npoints;
    w[i] = w[i] / sum;
  end;
  call WriteMember(selected,1,w);
  /* convert weights to x-coordinate form */
  x[1] = 0;
  x[2] = 0;
  do i = 1 to npoints;
    x[1] + w[i] * cvxhull[i,1];
    x[2] + w[i] * cvxhull[i,2];
  end;
end;

```

```

/* write out x coordinates to second segment */
call WriteMember(selected,2,x);
/* compute objective value */
r = (4 - 2.1*x[1]**2 + x[1]**4/3)*x[1]**2 + x[1]*x[2] +
    (-4 + 4*x[2]**2)*x[2]**2;
return(r);
endsub;
call SetObjFunc('sixhump',0);
array lower[1] /nosym;
array upper[1] /nosym;
call dynamic_array(lower, npoints);
call dynamic_array(upper, npoints);
do i = 1 to npoints;
    lower[i] = 0;
    upper[i] = 1;
end;
call SetBounds(lower, upper, 1);
array delta[3] /nosym (0.01 0.01 0.01);
call SetMut('delta', 'nchange', 1, 'delta', delta);
call SetMutProb(0.05);
call SetCross('Twopoint', 'alpha', 0.9);
call SetCrossProb(0.8);
call SetSel('tournament', 'size', 2);
call SetElite(3);
call Initialize('DEFAULT', 200);
call ContinueFor(200);
run;

```

Note that this problem uses the standard genetic operators and default initialization, even though they generate solutions that violate the constraints. This is possible because all solutions are passed into the user objective function for evaluation, where they are repaired to fit the constraints. The output is shown in Output 3.2.1.

Output 3.2.1 Nonlinear Objective with Constraints Using Repair Mechanism

PROC GA Optimum Values			
Objective			
-1.031617314			
Segment	Solution Element	Value	
1	1	0.2439660392	
1	2	0.5561415112	
1	3	0.1998924497	
2	1	-0.088147179	
2	2	0.712498123	

This objective function has a global minimum at -1.0316 , at two different points: $(x_1, x_2) = (-0.0898, 0.7126)$ and $(x_1, x_2) = (0.0898, -0.7126)$. The genetic algorithm can converge to either of

these minima, depending on the random number seed set by the SEED= option.

Example 3.3: Quadratic Objective with Linear Constraints, Using Bicriteria Approach

This example (Floudas and Pardalos 1992) illustrates the bicriteria approach to handling constraints. The problem has nine linear constraints and a quadratic objective function.

Minimize

$$f(x) = 5 \sum_{i=1}^4 x_i - 5 \sum_{i=1}^4 x_i^2 - \sum_{i=5}^{13} x_i$$

subject to

$$\begin{aligned} 2x_1 + 2x_2 + x_{10} + x_{11} &\leq 10 \\ 2x_1 + 2x_3 + x_{10} + x_{12} &\leq 10 \\ 2x_1 + 2x_3 + x_{11} + x_{12} &\leq 10 \\ -8x_1 + x_{10} &\leq 0 \\ -8x_2 + x_{11} &\leq 0 \\ -8x_3 + x_{12} &\leq 0 \\ -2x_4 - x_5 + x_{10} &\leq 0 \\ -2x_6 - x_7 + x_{11} &\leq 0 \\ -2x_8 - x_9 + x_{12} &\leq 0 \end{aligned}$$

and

$$\begin{aligned} 0 &\leq x_i \leq 1, & i &= 1, 2, \dots, 9 \\ 0 &\leq x_i \leq 100, & i &= 10, 11, 12 \\ 0 &\leq x_{13} \leq 1 \end{aligned}$$

In this example, the linear constraint coefficients are specified in the SAS data set lincon and passed to the GA procedure with the MATRIX1= option. The upper and lower bounds are specified in the bounds data set specified with a DATA1= option, which creates the array variables upper and lower, matching the variables in the data set.

```
/* Input linear constraint matrix */
data lincon;
    input A1-A13 b;
    datalines;
    2  2  0  0  0  0  0  0  0  1  1  0  0 10
    2  0  2  0  0  0  0  0  0  0  1  0  1 10
    2  0  2  0  0  0  0  0  0  0  0  1  1 10
   -8  0  0  0  0  0  0  0  0  0  1  0  0  0
    0 -8  0  0  0  0  0  0  0  0  0  1  0  0
    0  0 -8  0  0  0  0  0  0  0  0  1  0  0
    0  0  0 -2 -1  0  0  0  0  1  0  0  0  0
    0  0  0  0  0 -2 -1  0  0  0  1  0  0  0
    0  0  0  0  0  0  0 -2 -1  0  0  1  0  0
    ;
```

```

/* Input lower and upper bounds */
data bounds;
    input lower upper;
    datalines;
0 1
0 1
0 1
0 1
0 1
0 1
0 1
0 1
0 1
0 1
0 100
0 100
0 100
0 1
;

proc ga lastgen = out matrix1 = lincon
    seed = 12345 data1 = bounds;

```

Note also that the LASTGEN= option is used to designate a data set to store the final solution generation.

In the following statements, the solution encoding is specified, and a user function is defined and designated as the objective function.

```

call SetEncoding('R13R3');
nvar = 13;
ncon = 9;
function quad(selected[,], matrix1[,*], nvar, ncon);

    array x[1] /nosym;

    array r[3] /nosym;

    array violations[1] /nosym;

    call dynamic_array(x, nvar);

    call dynamic_array(violations, ncon);

    call ReadMember(selected, 1, x);

    sum1 = 0;

    do i = 1 to 4;

        sum1 + x[i] - x[i] * x[i];

    end;

    sum2 = 0;

```

```

do i = 5 to 13;

    sum2 + x[i];

end;

obj = 5 * sum1 - sum2;

call EvaluateLC(matrix1,violations,sumvio,selected,1);

r[1] = obj;

r[2] = sumvio;

call WriteMember(selected,2,r);

return(obj);

endsub;

call SetObjFunc('quad',0);

```

The [SetEncoding](#) call specifies two real-valued segments. The first segment, R13, holds the 13 variables, and the second segment, R3, holds the two objective criteria and the marker for Pareto optimality. As described in the section “[Defining an Objective Function](#)” on page 68, the first parameter of the objective function is a numeric array that designates which member of the solution population is to be evaluated. When the quad function is called by the GA procedure during the optimization process, the matrix1, nvar, and ncon parameters receive the values of the corresponding global variables; nvar is set to the number of variables, and ncon is set to the number of linear constraints. The function computes the original objective as the first objective criterion, and the magnitude of constraint violation as the second. With the first dynamic_array call, it allocates a working array, x, large enough to hold the number of variables, and a second array, violations, large enough to tabulate each constraint violation. The [ReadMember](#) call fills x with the elements of the first segment of the solution, and then the function computes the original objective $f(x)$. The [EvaluateLC](#) call is used to compute the linear constraint violation. The objective and sum of the constraint violations are then stored in the array r, and written back to the second segment of the solution with the [WriteMember](#) call. Note that the third element of r is not modified, because that element of the segment is used to store the Pareto-optimality mark, which cannot be determined until all the solutions have been evaluated.

Next, a user routine is defined and designated to be an update routine. This routine is called once at each iteration, after all the solutions have been evaluated with the quad function. The following program illustrates this:

```

subroutine update(popsiz);

    /* find pareto-optimal set */

    array minmax[3] /nosym (-1 -1 0);

    array results[1,1] /nosym;

    array scratch[1] /nosym;

```

```

call dynamic_array(scratch, popsize);

call dynamic_array(results,popsize,3);

/* read original and constraint objectives, stored in
   * solution segment 2, into array */

call GetSolutions(results,popsize,2);

/* mark the pareto-optimal set */

call MarkPareto(scratch, npareto,results, minmax);

/* transfer the results to the solution segment */

do i = 1 to popsize;

    results[i,3] = scratch[i];

end;

/* write updated segment 2 back into solution population */

call UpdateSolutions(results,popsize,2);

/* Set Elite parameter to preserve the first 15 pareto-optimal
   * solutions */

if npareto < 16 then

    call SetElite(npareto);

else

    call SetElite(15);

endsub;

call SetUpdateRoutine('update');

```

This subroutine has one parameter, `popsize`, defined within the GA procedure, which is expected to be the population size. The working arrays `results`, `scratch`, and `minmax` are declared. The `minmax` array is to be passed to a [MarkPareto call](#), and is initialized to specify that the first two elements (the original objective and constraint violation) are to be minimized and the third element is not to be considered. The `results` and `scratch` arrays are then dynamically allocated to the dimensions required by the population size.

Next, the `results` array is filled with the second segment of the solution population, with the [GetSolutions call](#). The `minmax` and `results` arrays are passed as inputs to the [MarkPareto call](#), which returns the number of Pareto optimal solutions in the `npareto` variable. The `MarkPareto` call also sets the elements of the `scratch` array to 1 if the corresponding solution is Pareto-optimal, and to 0 otherwise. The next loop then records the results in the `scratch` array in the third column of the `results` array, effectively marking the Pareto-optimal

solutions. The updated solution segments are written back to the population with the UpdateSolutions call.

The final step in the update routine is to set the elite selection parameter to guarantee the survival of at least a minimum of 15 of the fittest (Pareto-optimal) solutions through the selection process.

With the following statements, a routine is defined and designated as a fitness comparison routine with a [SetCompareRoutine](#) call. This routine works in combination with the update routine to evolve the solution population toward Pareto optimality and constraint satisfaction.

```
function paretocomp(selected[*]);

    array member1[3] /nosym;

    array member2[3] /nosym;

    call ReadCompare(selected,2,1, member1);

    call ReadCompare(selected,2,2, member2);

    /* if one member is in the pareto-optimal set
       * and the other is not, then it is the
       * most fit */

    if(member1[3] > member2[3]) then

        return(1);

    if(member2[3] > member1[3]) then

        return(-1);

    /* if both are in the pareto-optimal set, then
       * the one with the lowest constraint violation
       * is the most fit */

    if(member1[3] = 1) then do;

        if member1[2] <= member2[2] then

            return(1);

        return( -1);

    end;

    /* if neither is in the pareto-optimal set, then
       * take the one that dominates the other */
```

```

    if (member1[2] <= member2[2]) &
        (member1[1] <= member2[1]) then
        return(1);

    if (member2[2] <= member1[2]) &
        (member2[1] <= member1[1]) then
        return(-1);

    /* if neither dominates, then consider fitness to be
       * the same */
    return( 0);

endsub;

call SetSel('tournament', 'size', 2);

call SetCompareRoutine('paretocomp');

```

The **PARETOCOMP** subroutine is called in the selection process to compare the fitness of two competing solutions. The first parameter, `selected`, designates the two solutions to be compared.

The `ReadCompare` calls retrieve the second segments of the two solutions, where the objective criteria are stored, and writes the segments into the `member1` and `member2` arrays. The logic that follows first checks for the case where only one solution is Pareto optimal, and returns it. If both the solutions are Pareto optimal, then the one with the smallest constraint violation is chosen. If neither solution is Pareto optimal, then the dominant solution is chosen, if one exists. If neither solution is dominant, then no preference is indicated. After the function is defined, it is designated as a fitness comparison routine with the [SetCompareRoutine](#) call.

Next, subroutines are defined and designated as user crossover and mutation operators:

```

/* set up crossover parameters */

subroutine Cross1(selected[*], alpha);

    call Cross(selected,1,'twopoint', alpha);

endsub;

call SetCrossRoutine('Cross1',2,2);

alpha = 0.5;

call SetCrossProb(0.8);

/* set up mutation parameters */

```



```

subroutine Mut1(selected[*], delta[*]);

    call Mutate(selected,1,'delta',delta,1);

endsub;

call SetMutRoutine('Mut1');

array delta[13] /nosym (.5 .5 .5 .5 .5 .5 .5 .5 .5 10 10 10 .1);

call SetMutProb(0.05);

```

These routines execute the standard genetic operators *twopoint* for crossover and *delta* for mutation; see the section “[Using Standard Genetic Operators and Objective Functions](#)” on page 54 for a description of each. The alpha and delta variables defined in the procedure are passed as parameters to the user operators, and the crossover and mutation probabilities are set with the [SetCrossProb](#) and [SetMutProb](#) calls.

At this point, the GA procedure is directed to initialize the first population and begin the optimization process:

```

/* Initialize first population */

call SetBounds(lower, upper);

popsize = 100;

call Initialize('DEFAULT',popsize);

call ContinueFor(500);

run;

```

First, the upper and lower bounds are established with values in the lower and upper array variables, which were set up by the DATA1= option in the PROC GA statement. The [SetBounds](#) call sets the bounds for the first segment, which is the default if none is specified in the call. The desired population size of 100 is stored in the popsize variable, so it will be passed to the update subroutine as the popsize parameter. The [Initialize](#) call specifies the default initialization, which generates values randomly distributed between the lower and upper bounds for the first encoding segment. Since no bounds were specified for the second segment, it is filled with zeros. The [ContinueFor](#) call sets the requested number of iterations to 500, and the RUN statement ends the GA procedure input and begins the optimization process. The output of the procedure is shown in [Output 3.3.1](#).

Output 3.3.1 Bicriteria Constraint Handling Example Output

Bicriteria Constraint Handling Example	
PROC GA Optimum Values	
Objective	
-14.99871988	

Output 3.3.1 *continued*

Segment	Solution	
	Element	Value
1	1	1
1	2	0.9999997423
1	3	0.9999991741
1	4	0.9999997454
1	5	0.9999982195
1	6	1
1	7	0.9999674484
1	8	0.9999961238
1	9	0.9999691145
1	10	2.9999914332
1	11	2.9999103023
1	12	2.9988939259
1	13	1
2	1	-14.99871988
2	2	0
2	3	1

The minimum value of $f(x)$ is -15 at $x^* = (1, 1, 1, 1, 1, 1, 1, 1, 1, 3, 3, 3, 1)$.

References

- Floudas, C. A. and Pardalos, P. M. (1992), *Recent Advances in Global Optimization*, Princeton, NJ: Princeton University Press.
- Michalewicz, Z. (1996), *Genetic Algorithms + Data Structures = Evolution Programs*, New York: Springer-Verlag.

Subject Index

- bicriteria comparison strategy, 73
- choosing problem encoding, 21
- constraints, 72
 - bicriteria comparison strategy, 73
 - direct comparison strategy, 73
 - encoding strategy, 72
 - penalty strategy, 73
 - repair strategy, 72
- controlling selection process, 22
- creating initial generation, 24
- crossover operators, 55
- defining
 - fitness comparison routine, 63
 - genetic operators, 64
 - initialization routine, 69
 - objective functions, 68
 - update routine, 67
- direct comparison strategy, 73
- duel selector, 71
- encoding strategy, 72
- fitness comparison routine
 - defining, 63
- GA procedure
 - debugging options, 39
 - overview, 16
 - program statements, 38
- genetic algorithms, 16
- genetic operators, 54
 - defining, 64
- heuristics, 72
- initialization routine
 - defining, 69
- initializing problem data, 19
- local optimizations, 72
- monitoring progress, 25
- multiple objectives, 73
- multisegment encoding, 53
- mutation operators, 61
- objective functions, 54, 63
 - defining, 68
- online documentation, 11
- overview
 - GA procedure, 16
- penalty strategy, 73
- program statements
 - GA procedure, 38
- repair strategy, 72
- reporting results, 25
- selection strategy, 71
- setting crossover parameters, 23
- setting mutation parameters, 24
- setting objective function, 22
- tournament selector, 71
- update routine
 - defining, 67

Syntax Index

ABORT statement
GA program statements, 39

DATAN= option
PROC GA statement, 29

DO statement
GA program statements, 39

FIRSTGEN= option
PROC GA statement, 29

GA procedure
ContinueFor Call, 30
Cross call, 31
dynamic_array, 31
EvaluateLC, 32
GetDimensions, 33
GetObjValues, 33
GetSolutions Call, 33
Initialize Call, 34
MarkPareto Call, 35
Mutate Call, 36
Objective Call, 37
PackBits Call, 38
PROC GA statement, 28
ReadChild Call, 40
ReadCompare Call, 41
ReadMember Call, 41
ReadParent Call, 42
ReEvaluate Call, 42
SetBounds Call, 42
SetCross Call, 43
SetCrossProb Call, 44
SetCrossRoutine Call, 45
SetElite Call, 45
SetEncoding Call, 45
SetFinalize Call, 46
SetMut Call, 46
SetMutProb Call, 47
SetMutRoutine Call, 43, 47
SetObj Call, 48
SetObjFunc Call, 48
SetProperty Call, 49
SetSel Call, 49
SetUpdateRoutine Call, 50
ShellSort Call, 50
Shuffle Call, 51
UnpackBits Function, 51
UpdateSolutions Call, 52

WriteChild Call, 52
WriteMember Call, 52

LASTGEN= option
PROC GA statement, 29
LIBRARY= option
PROC GA statement, 29

MATRIXn= option
PROC GA statement, 30

NITER= option
PROC GA statement, 30
NOVALIDATE= option
PROC GA statement, 30
NOVALIDATEWARNING= option
PROC GA statement, 30

OTHERWISE statement
GA program statements, 40

PROC GA statement, *see* GA procedure
statement options, 28

PUT statement
GA program statements, 39

SEED= option
PROC GA statement, 30

SELECT statement
GA program statements, 40

WHEN statement
GA program statements, 40

Your Turn

We welcome your feedback.

- If you have comments about this book, please send them to **`yourturn@sas.com`**. Include the full title and page numbers (if applicable).
- If you have comments about the software, please send them to **`suggest@sas.com`**.

SAS® Publishing Delivers!

Whether you are new to the work force or an experienced professional, you need to distinguish yourself in this rapidly changing and competitive job market. SAS® Publishing provides you with a wide range of resources to help you set yourself apart. Visit us online at support.sas.com/bookstore.

SAS® Press

Need to learn the basics? Struggling with a programming problem? You'll find the expert answers that you need in example-rich books from SAS Press. Written by experienced SAS professionals from around the world, SAS Press books deliver real-world insights on a broad range of topics for all skill levels.

support.sas.com/saspress

SAS® Documentation

To successfully implement applications using SAS software, companies in every industry and on every continent all turn to the one source for accurate, timely, and reliable information: SAS documentation. We currently produce the following types of reference documentation to improve your work experience:

- Online help that is built into the software.
- Tutorials that are integrated into the product.
- Reference documentation delivered in HTML and PDF – **free** on the Web.
- Hard-copy books.

support.sas.com/publishing

SAS® Publishing News

Subscribe to SAS Publishing News to receive up-to-date information about all new SAS titles, author podcasts, and new Web site features via e-mail. Complete instructions on how to subscribe, as well as access to past issues, are available at our Web site.

support.sas.com/spn



**THE
POWER
TO KNOW®**

