SAS/OR® 9.22 User's Guide
Constraint Programming

Contents

Credits

Documentation

Writing | Gehan A. Corea

Editing | Edward Huddleston, Anne Jones

Documentation Support | Tim Arnold, Natalie Baerlocher, Liping Cai, Richard Liu, Jianzhe Luo, Lindsey Puryear

Technical Review | Tonya Chapman, Edward P. Hughes, Charles B. Kelly, Radhika V. Kulkarni, Rob Pratt

---

Software

PROC CLP | Gehan A. Corea, Keqi Yan, Tien-yi D. Shaw

Support Groups

Software Testing | Liping Cai, Charles B. Kelly, Rob Pratt, Lindsey Puryear

Technical Support | Tonya Chapman

Chapter 1

What's New in SAS/OR 9.22

Contents

Overview

SAS/OR 9.22 continues the improvements that were delivered starting with SAS/OR 9.2. Several new and enhanced features expand the scale and scope of problems that SAS/OR software can address. These enhancements also make it easier for you to use the SAS/OR capabilities. Brief descriptions of these new features are presented in the following sections.

Highlights of Enhancements in SAS/OR 9.22

Highlights of the changes include the following:

- You can customize the format of the time axis on the Gantt chart.

- You can import and convert Microsoft Project data that has been saved in XML format.

- The CLP procedure is now production with the exception of the scheduling related constraints.

- The OPTMODEL procedure supports named problems to enable easy manipulation of multiple subproblems.

- The IPNLP and NLPU solvers support new techniques for large-scale optimization.

- SAS Simulation Studio 1.5 is a new graphical application for discrete event simulation and is included with SAS/OR software. Documentation is available at the following link: http://support.sas.com/documentation/onlinedoc/simstudio/index.html

More information about the changes and enhancements is provided in this chapter. Details can be found in the relevant volumes of the *SAS/OR 9.22 User's Guide* and in the *SAS Simulation Studio 1.5 User's Guide*.

Highlights of Enhancements in SAS/OR 9.2

Some users are moving directly from SAS/OR 9.1.3 to SAS/OR 9.22. The following are some of the major enhancements that were introduced in SAS/OR 9.2:

- The MPSOUT= option directs procedures to save input problem data in an MPS-format SAS data set. The MPSOUT= option is available in the LP, NETFLOW, and OPTLP procedures.

- The IIS= option for the LP solver enables you to identify, for an infeasible linear program, constraints and variable bounds that form an irreducible infeasible set (IIS). The IIS= option is available in the OPTLP and OPTMODEL procedures.

- The value "2" for the PRINTLEVEL= option directs procedures to produce an ODS table called "ProblemStatistics" in addition to the "ProblemSummary" and "SolutionSummary" ODS tables that are produced for PRINTLEVEL=1. The PRINTLEVEL=2 option is available in the INTPOINT, OPTLP, and OPTMILP procedures.

- The %SASTOMSP macro converts data sets that are used by the CPM and PM procedures into an MDB file that is readable by Microsoft Project.

- Several call routines in the GA procedure were replaced by new call routines.

- The CLP procedure features improved algorithms for the "all-different" constraint in addition to several extensions to the edge-finder algorithm for resource-constrained scheduling.

For more information, see support.sas.com/whatsnewor92.

SAS/OR Documentation

SAS/OR software is documented in the following volumes:

- *SAS/OR User's Guide: Bills of Material Processing*

- *SAS/OR User's Guide: Constraint Programming*

- *SAS/OR User's Guide: Local Search Optimization*

- *SAS/OR User's Guide: Mathematical Programming*

- *SAS/OR User's Guide: Project Management*

- *SAS/OR User's Guide: QSIM Application*

- *SAS Simulation Studio 1.5: User's Guide*

Online help can also be found under the corresponding classification.

The GANTT Procedure

The GANTT procedure produces a Gantt chart, which is a graphical tool for representing schedule-related information. PROC GANTT provides support for displaying multiple schedules, precedence relationships, calendar information, milestones, reference lines, labeling, and so on. New in SAS/OR 9.22 is the TIMEAXISFORMAT= option in the CHART statement which provides the capability to customize the format of the time axis on the Gantt chart for up to three rows. Each row can be formulated using a predefined SAS format or a user-defined format.

Microsoft Project Conversion Macros

The SAS macro %MSPTOSAS converts Microsoft Project 98 (and later) data into SAS data sets that can be used as input for project scheduling with SAS/OR software. This macro generates the necessary SAS data sets, determines the values of the relevant options, and invokes the SAS/OR PM procedure with the converted project data. The %MSPTOSAS macro enables you to use Microsoft Project for the input of project data and still take advantage of the excellent SAS/OR project and resource scheduling capabilities. New in SAS/OR 9.22 is the capability to import and convert Microsoft Project data that has been saved in XML format. This feature is experimental.

The experimental %SASTOMSP macro converts data sets that are used by the CPM and PM procedures into a Microsoft Access Database (MDB) file that is readable by Microsoft Project. The macro converts information that is common to PROC CPM, PROC PM, and Microsoft Project; this information includes hierarchical relationships, precedence relationships, time constraints, resource availabilities, resource requirements, project calendars, resource calendars, task calendars, holiday information, and work-shift information. In addition, the early and late schedules, the actual start and finish times, the resource-constrained schedule, and the baseline schedule are also extracted and stored as start-finish variables.

Execution of the %MSPTOSAS and %SASTOMSP macros requires SAS/ACCESS® software.

The CLP Procedure

The CLP procedure is a finite-domain constraint programming solver for solving constraint satisfaction problems (CSPs) with linear, logical, global, and scheduling constraints. The CLP procedure is production in SAS/OR 9.22 with the exception of the scheduling-related constraints.

New in SAS/OR 9.22 are the GCC and ELEMENT statements for defining global cardinality constraints (GCC) and element constraints, respectively. The GCC statement enables you to bound the number of times that a specific value gets assigned to a set of variables. The ELEMENT statement enables you to define dependencies, not necessarily functional, between variables and to define noncontiguous domains.

The USECONDATAVARS= option enables you to implicitly define numeric variables in the CONDATA= data set. The TIMETYPE= option enables you to set the units (real time or CPU time) of the MAXTIME= parameter. The _ORCLP_ macro variable has been enhanced to provide more information about procedure status and solution status.

There are also several changes and enhancements to the scheduling capabilities in SAS/OR 9.22. Support for multiple-capacity resources has been added in the RESOURCE statement and the Activity data set. The REQUIRES statement syntax for specifying multiple resource requirements has changed. The format of the Activity data set has changed to a more compact form with a fixed number of variables. A new Resource data set, specified with the RESDATA= option, enables you to define resources, resource pools, and resource attributes in compact form. The format of the Schedule data set has been enhanced to separate time and schedule related observations. Two new schedule-related output data sets, SCHEDTIME= and SCHEDRES=, have been added; they contain time assignment and resource assignment information, respectively.

The OPTMODEL Procedure

The OPTMODEL procedure provides a modeling environment that is tailored to building, solving, and maintaining optimization models. This makes the process of translating the symbolic formulation of an optimization model into PROC OPTMODEL virtually transparent, because the modeling language mimics the symbolic algebra of the formulation as closely as possible. PROC OPTMODEL also streamlines and simplifies the critical process of populating optimization models with data from SAS data sets. All of this transparency produces models that are more easily inspected for completeness and correctness, more easily corrected, and more easily modified, whether through structural changes or through the substitution of new data for old data.

The OPTMODEL procedure consists of the powerful OPTMODEL modeling language and access to state-of-the-art solvers for several classes of mathematical programming problems.

Seven solvers are available to OPTMODEL, as listed in Table 1.1.

Table 1.1 List of OPTMODEL Solvers

Problem	Solver
Linear programming	**LP**
Mixed integer programming	**MILP**
Quadratic programming (Experimental)	**QP**
Nonlinear programming, unconstrained	**NLPU**
General nonlinear programming	**NLPC**
General nonlinear programming	**SQP**
General nonlinear programming	**IPNLP**

In SAS/OR 9.22, the OPTMODEL procedure adds several new features. First, PROC OPTMODEL supports named problems to enable easy manipulation of multiple subproblems. The PROBLEM declaration declares a named problem and the USE PROBLEM statement makes it active. Objectives can now be declared as arrays, so they can provide separate objectives for arrays of named problems.

Implicit variables, created via the IMPVAR declaration, allow optimization expressions to be referred to by name in a model. Implicit variables can be evaluated more efficiently than by repeating the same complex expression in multiple places.

Problem components can be accessed with aliases such as _VAR_ and _CON_, which respectively aggregate all of the variables and constraints in a problem. This allows convenient processing of all of the problem components of a given kind for printing, model expansion, and other purposes. The new suffixes .NAME and .LABEL can be used to track the identity of problem components.

Function and subroutine calls can use the "OF array-name[*]" syntax to pass an OPTMODEL array to a called routine for uses such as sorting.

The NUMBER, STRING, and SET declarations allow initial values for arrays to be supplied using an INIT clause with a list of initialization values.

The SOLVE statement supports the RELAXINT keyword to solve a problem while temporarily relaxing the integrality restriction on variables.

Analytic derivatives are now generated for most SAS library functions. The OPTMODEL procedure can use threading on systems with multiple processors to speed up evaluation of nonlinear Hessian models.

Starting with SAS/OR 9.22, the IPNLP and NLPU solvers support new techniques for large-scale optimization. The nonlinear solver IPNLP has been equipped with two new techniques. The first technique, TECH=IPKRYLOV, is appropriate for large-scale nonlinear optimization problems that can contain many thousands of variables or constraints or both. It uses exact second derivatives to calculate the search directions. Its convergence is achieved by using a trust-region framework that guides the algorithm towards the solution of the optimization problem. The second technique, TECH=IPQN, uses a quasi-Newton method and line-search framework to solve the optimization problem. As such it needs to calculate only the first derivatives of the objective and constraints. This method is more appropriate for problems where the second derivatives of the objective and constraints either are not available or are expensive to compute.

The unconstrained solver NLPU has been equipped with a new technique called TECH=CGTR. This technique uses the conjugate gradient method to solve large-scale unconstrained and bound constrained optimization problems.

The OPTMILP Procedure

The OPTMILP procedure solves mixed-integer linear programming problems with a linear-programming-based branch-and-bound algorithm that has been improved for SAS/OR 9.22. The algorithmic improvements result from incorporating new techniques in the presolver and cutting planes, better application of primal heuristics, an improved branch-and-bound strategy, and an improved strategy for handling feasibility problems. Improvements to the presolver include variable and constraint reductions based on logical implications among binary variables and generalized variable substitutions. Two new cutting plane routines (mixed 0-1 lifted inequalities and zero-half cuts) have been added, and improvements have been made to clique, Gomory mixed integer, and mixed integer rounding (MIR) cutting plane routines.

The resulting improvements in efficiency enable you to use PROC OPTMILP to solve larger and more complex optimization problems in a shorter time than with previous SAS/OR releases.

SAS Simulation Studio

SAS Simulation Studio is a discrete event simulation application for modeling the operation of call centers, supply chains, emergency rooms, and other real-world systems in which there are significant random elements (timing and length of events, requirements, and so on). Its graphical user interface provides a full set of tools and components for building, executing, and analyzing the data that are generated by discrete event simulation models. SAS Simulation Studio provides extensive modeling and analysis tools suitable for both novice and advanced simulation users.

SAS Simulation Studio integrates fully with JMP software to provide experimental design capabilities for evaluating and analyzing your simulation models. Any of the JMP and SAS statistical analysis tools can be used, either to analyze results after the simulation model is run or to perform embedded analyses that occur while the simulation model is running.

SAS Simulation Studio 1.5 has been included with SAS/OR software since its release in August 2009.

Chapter 2

Using This Book

Contents

Purpose

SAS/OR User's Guide: Constraint Programming provides a complete reference for the constraint programming procedures in SAS/OR software. This book serves as the primary documentation for the CLP procedure.

"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 10 refers to other documents that contain related information.

Organization

Chapter 3 describes the CLP 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 most procedures. The statements and syntax necessary to run each procedure are presented in a uniform format throughout this book.

The following list summarizes the types of information provided for each 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. Following these tables, the PROC statement is described, and then all other statements are described in alphabetical order.

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: Constraint Programming* 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.
oblique	is used to indicate an option variable for which you must supply a value (for example, DUPLICATE= *dup* indicates that you must supply a value for *dup*).
italic	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**. (Alternatively, you can type **help OR** in the command line.) 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: Constraint Programming,* you may find these other 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: Local Search Optimization
> provides documentation for the local search optimization procedure in SAS/OR software. This book serves as the primary documentation for the GA procedure, which uses genetic algorithms to solve optimization problems.

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 INTPOINT, LP, NETFLOW, and NLP procedures, in addition to the newer 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: 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, as well as the PROJMAN Application, a graphical user interface for project management.

SAS/OR User's Guide: The QSIM Application
> provides documentation for the QSIM application, which is used to build and analyze models of queueing systems using discrete event simulation. This book shows you how to build models using the simple point-and-click graphical user interface, how to run the models, and how to collect and analyze the sample data to give you insight into the behavior of the system.

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 1.5: 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 CLP Procedure

Contents

Overview: CLP Procedure

The CLP procedure is a finite-domain constraint programming solver for constraint satisfaction problems (CSPs) with linear, logical, global, and scheduling constraints. In addition to having an expressive syntax for representing CSPs, the CLP procedure features powerful built-in consistency routines and constraint propagation algorithms, a choice of nondeterministic search strategies, and controls for guiding the search mechanism that enable you to solve a diverse array of combinatorial problems.

The scheduling-related constraints are experimental in the current version of the CLP procedure.

The Constraint Satisfaction Problem

Many important problems in areas such as artificial intelligence (AI) and operations research (OR) can be formulated as constraint satisfaction problems. A CSP is defined by a finite set of variables that take values from finite domains and a finite set of constraints that restrict the values that the variables can simultaneously take.

More formally, a CSP can be defined as a triple $\langle X, D, C \rangle$:

- $X = \{x_1, \ldots, x_n\}$ is a finite set of *variables*.

- $D = \{D_1, \ldots, D_n\}$ is a finite set of *domains*, where D_i is a finite set of possible values that the variable x_i can take. D_i is known as the *domain* of variable x_i.

- $C = \{c_1, \ldots, c_m\}$ is a finite set of *constraints* that restrict the values that the variables can simultaneously take.

The domains need not represent consecutive integers. For example, the domain of a variable could be the set of all *even* numbers in the interval [0, 100]. A domain does not even need to be totally numeric. In fact, in a scheduling problem with resources, the values are typically multidimensional. For example, an activity can be considered as a variable, and each element of the domain would be an *n-tuple* that represents a start time for the activity and one or more resources that must be assigned to the activity that corresponds to the start time.

A solution to a CSP is an assignment of values to the variables in order to satisfy all the constraints. The problem amounts to finding one or more solutions, or possibly determining that a solution does not exist.

The CLP procedure can be used to find one or more (and in some instances, all) solutions to a CSP with linear, logical, global, and scheduling constraints. The numeric components of all variable domains are assumed to be integers.

Techniques for Solving CSPs

Several techniques for solving CSPs are available. Kumar (1992) and Tsang (1993) present a good overview of these techniques. It should be noted that the satisfiability problem (SAT) (Garey and Johnson 1979) can be regarded as a CSP. Consequently, most problems in this class are NP-complete problems, and a backtracking search mechanism is an important technique for solving them (Floyd 1967).

One of the most popular tree search mechanisms is chronological backtracking. However, a chronological backtracking approach is not very efficient due to the late detection of conflicts; that is, it is oriented toward *recovering* from failures rather than *avoiding* them to begin with. The search space is reduced only after detection of a failure, and the performance of this technique is drastically reduced with increasing problem size. Another drawback of using chronological backtracking is encountering repeated failures due to the same reason, sometimes referred to as "thrashing." The presence of late detection and "thrashing" has led researchers to develop consistency techniques that can achieve superior pruning of the search tree. This strategy employs an active use, rather than a passive use, of constraints.

Constraint Propagation

A more efficient technique than backtracking is that of constraint propagation, which uses consistency techniques to effectively prune the domains of variables. Consistency techniques are based on the idea of *a priori* pruning, which uses the constraint to reduce the domains of the variables. Consistency techniques are also known as relaxation algorithms (Tsang 1993), and the process is also referred to as problem reduction, domain filtering, or pruning.

One of the earliest applications of consistency techniques was in the AI field in solving the scene labeling problem, which required recognizing objects in three-dimensional space by interpreting two-dimensional line drawings of the object. The Waltz filtering algorithm (Waltz 1975) analyzes line drawings by systematically labeling the edges and junctions while maintaining consistency between the labels.

An effective consistency technique for handling resource capacity constraints is edge finding (Applegate and Cook 1991). Edge-finding techniques reason about the processing order of a set of activities that require a given resource or set of resources. Some of the earliest work related to edge finding can be attributed to Carlier and Pinson (1989), who successfully solved MT10, a well-known 10×10 job shop problem that had remain unsolved for over 20 years (Muth and Thompson 1963).

Constraint propagation is characterized by the extent of propagation (also referred to as the level of consistency) and the domain pruning scheme that is followed: domain propagation or interval propagation. In practice, interval propagation is preferred over domain propagation because of its lower computational costs. This mechanism is discussed in detail in Van Hentenryck (1989). However, constraint propagation is not a complete solution technique and needs to be complemented by a search technique in order to ensure success (Kumar 1992).

Finite-Domain Constraint Programming

Finite-domain constraint programming is an effective and complete solution technique that embeds incomplete constraint propagation techniques into a nondeterministic backtracking search mechanism, implemented as follows. Whenever a node is visited, constraint propagation is carried out to attain a desired level of consistency. If the domain of each variable reduces to a singleton set, the node represents a solution to the CSP. If the domain of a variable becomes empty, the node is pruned. Otherwise a variable is selected, its domain is distributed, and a new set of CSPs is generated, each of which is a child node of the current node. Several factors play a role in determining the outcome of this mechanism, such as the extent of propagation (or level of consistency enforced), the variable selection strategy, and the variable assignment or domain distribution strategy.

For example, the lack of any propagation reduces this technique to a simple generate-and-test, whereas performing consistency on variables already selected reduces this to chronological backtracking, one of the systematic search techniques. These are also known as look-back schemas, because they share the disadvantage of late conflict detection. Look-ahead schemas, on the other hand, work to prevent future conflicts. Some popular examples of look-ahead strategies, in increasing degree of consistency level, are forward checking (FC), partial look ahead (PLA), and full look ahead (LA) (Kumar 1992). Forward checking enforces consistency between the current variable and future variables; PLA and LA extend this even further to pairs of not yet instantiated variables.

Two important consequences of this technique are that inconsistencies are discovered early and that the current set of alternatives that are coherent with the existing partial solution is dynamically maintained. These consequences are powerful enough to prune large parts of the search tree, thereby reducing the "combinatorial explosion" of the search process. However, although constraint propagation at each node results in fewer nodes in the search tree, the processing at each node is more expensive. The ideal scenario is to strike a balance between the extent of propagation and the subsequent computation cost.

Variable selection is another strategy that can affect the solution process. The order in which variables are chosen for instantiation can have a substantial impact on the complexity of the backtrack search. Several heuristics have been developed and analyzed for selecting variable ordering. One of the more common ones is a dynamic heuristic based on the *fail first* principle (Haralick and Elliot 1980), which selects the variable whose domain has minimal size. Subsequent analysis of this heuristic by several researchers has validated this technique as providing substantial improvement for a significant class of problems. Another popular technique is to instantiate the most constrained variable first. Both these strategies are based on the principle of selecting the variable most likely to fail and to detect such failures as early as possible.

The domain distribution strategy for a selected variable is yet another area that can influence the performance of a backtracking search. However, good value-ordering heuristics are expected to be very problem-specific (Kumar 1992).

The CLP Procedure

The CLP procedure is a finite-domain constraint programming solver for CSPs. In the context of the CLP procedure, CSPs can be classified into the following two types which are determined by specification of the relevant output data set:

- A *standard CSP* is characterized by integer variables, linear constraints, array-type constraints, global constraints, and reify constraints. In other words, X is a finite set of integer variables, and C can contain linear, array, global, or logical constraints. Specifying the OUT= option in the PROC CLP statement indicates to the CLP procedure that the CSP is a standard type CSP. As such, the procedure expects only VARIABLE, ALLDIFF, ELEMENT, GCC, LINCON, REIFY, ARRAY, and FOREACH statements. You can also specify a Constraint data set by using the CONDATA= option in the PROC CLP statement in lieu of, or in combination with, VARIABLE and LINCON statements.

- A *scheduling CSP* is characterized by activities, temporal constraints, and resource requirement constraints. In other words, X is a finite set of activities, and C is a set of temporal constraints and resource requirement constraints. Specifying one of the SCHEDULE=, SCHEDRES=, or SCHEDTIME= options in the PROC CLP statement indicates to the CLP procedure that the CSP is a scheduling type CSP. As such, the procedure expects only ACTIVITY, RESOURCE, REQUIRES, and SCHEDULE statements. You can also specify an Activity data set by using the ACTDATA= option in the PROC CLP statement in lieu of, or in combination with, the ACTIVITY, RESOURCE, and REQUIRES statements. Activities can be defined using the Activity data set or the ACTIVITY statement. Precedence relationships between activities must be defined using the ACTDATA= data set. Resource requirements of activities can be defined using the Activity data set or the RESOURCE and REQUIRES statements.

The output data sets contain any solutions determined by the CLP procedure. For more information about the format and layout of the output data sets, see "Solution Data Set" on page 44 and "Schedule Data Set" on page 49.

Consistency Techniques

The CLP procedure features a full look-ahead algorithm for standard CSPs that follows a strategy of maintaining a version of generalized arc consistency that is based on the AC-3 consistency routine (Mackworth 1977). This strategy maintains consistency between the selected variables and the unassigned variables and also maintains consistency between unassigned variables. For the scheduling CSPs, the CLP procedure uses a forward-checking algorithm, an arc-consistency routine for maintaining consistency between unassigned activities, and energetic-based reasoning methods for resource-constrained scheduling that feature the edge-finder algorithm (Applegate and Cook 1991). You can elect to turn off some of these consistency techniques in the interest of performance.

Selection Strategy

A search algorithm for CSPs searches systematically through the possible assignments of values to variables. The order in which a variable is selected can be based on a *static* ordering, which is determined before the search begins, or on a *dynamic* ordering, in which the choice of the next variable depends on the current state of the search. The VARSELECT= option in the PROC CLP statement defines the variable selection strategy for a standard CSP. The default strategy is the dynamic MINR strategy, which selects the variable with the smallest range. The ACTSELECT= option in the SCHEDULE statement defines the activity selection strategy for a scheduling CSP. The default strategy is the RAND strategy, which selects an activity at random from the set of activities that begin prior to the earliest early finish time. This strategy was proposed by Nuijten (1994).

Assignment Strategy

After a variable or an activity has been selected, the assignment strategy dictates the value that is assigned to it. For variables, the assignment strategy is specified with the VARASSIGN= option in the PROC CLP statement. The default assignment strategy selects the minimum value from the domain of the selected variable. For activities, the assignment strategy is specified with the ACTASSIGN= option in the SCHEDULE statement. The default strategy of RAND assigns the time to the earliest start time, and the resources are chosen randomly from the set of resource assignments that support the selected start time.

Getting Started: CLP Procedure

The following examples illustrate the use of the CLP procedure in the formulation and solution of two well-known logical puzzles in the constraint programming community.

Send More Money

The Send More Money problem consists of finding unique digits for the letters D, E, M, N, O, R, S, and Y such that S and M are different from zero (no leading zeros) and the following equation is satisfied:

$$S\ E\ N\ D$$

$$+\ M\ O\ R\ E$$

$$\overline{}$$

$$M\ O\ N\ E\ Y$$

You can use the CLP procedure to formulate this problem as a CSP by representing each of the letters in the expression with an integer variable. The domain of each variable is the set of digits 0 through 9. The VARIABLE statement identifies the variables in the problem. The DOM= option defines the default domain for all the variables to be [0,9]. The OUT= option identifies the CSP as a standard type. The LINCON statement is used to define the linear constraint SEND + MORE = MONEY, and the restrictions that S and M cannot take the value zero. (Alternatively, you can simply specify the domain for S and M as [1,9] in the VARIABLE statement.) Finally, the ALLDIFF statement is specified to enforce the condition that the assignment of digits should be unique. The complete representation, using the CLP procedure, is as follows:

```
proc clp dom=[0,9]                   /* Define the default domain */
         out=out;                    /* Name the output data set  */
   var S E N D M O R E M O N E Y;    /* Declare the variables      */
   lincon                            /* Linear constraints         */
                                     /* SEND + MORE = MONEY        */
      1000*S + 100*E + 10*N + D + 1000*M + 100*O + 10*R + E
      =
      10000*M + 1000*O + 100*N  + 10*E + Y,
      S<>0,                          /* No leading zeros           */
      M<>0;
   alldiff();     /* All variables have pairwise distinct values*/
run;
```

The solution data set produced by the CLP procedure is shown in Figure 3.1.

Figure 3.1 Solution to SEND + MORE = MONEY

S	E	N	D	M	O	R	Y
9	5	6	7	1	0	8	2

The unique solution to the problem determined by the CLP procedure is as follows:

$$9\,5\,6\,7$$

$$+\,1\,0\,8\,5$$

$$\overline{}$$

$$1\,0\,6\,5\,2$$

Eight Queens

The Eight Queens problem is a special instance of the N-Queens problem, where the objective is to position N queens on an $N{\times}N$ chessboard such that no two queens attack each other. The CLP procedure provides an expressive constraint for variable arrays that can be used for solving this problem very efficiently.

You can model this problem by using a variable array A of dimension N, where $A[i]$ is the row number of the queen in column i. Since no two queens can be in the same row, it follows that all the $A[i]$'s must be pairwise distinct.

In order to ensure that no two queens can be on the same diagonal, the following should be true for all i and j:

$$A[j] - A[i] <> j - i$$

and

$$A[j] - A[i] <> i - j$$

In other words,

$$A[i] - i <> A[j] - j$$

and

$$A[i] + i <> A[j] + j$$

Hence, the $(A[i] + i)$'s are pairwise distinct, and the $(A[i] - i)$'s are pairwise distinct.

These two conditions, in addition to the one requiring that the $A[i]$'s be pairwise distinct, can be formulated using the FOREACH statement.

One possible such CLP formulation is presented as follows:

```
proc clp out=out
        varselect=fifo; /* Variable Selection Strategy          */
   array A[8] (A1-A8);   /* Define the array A                   */
   var (A1-A8)=[1,8];    /* Define each of the variables in the array */
                         /* Initialize domains                   */
   /* A[i] is the row number of the queen in column i*/
   foreach(A, DIFF,  0); /* A[i] 's are pairwise distinct */
   foreach(A, DIFF, -1); /* A[i] - i 's are pairwise distinct */
   foreach(A, DIFF,  1); /* A[i] + i 's are pairwise distinct */
run;
```

The ARRAY statement is required when you are using a FOREACH statement, and it defines the array A in terms of the eight variables A1–A8. The domain of each of these variables is explicitly specified in the VARIABLE statement to be the digits 1 through 8 since they represent the row number on an 8×8 board. FOREACH(A, DIFF, 0) represents the constraint that the $A[i]$'s are different. FOREACH(A, DIFF, –1) represents the constraint that the $(A[i] - i)$'s are different, and FOREACH(A, DIFF, 1) represents the constraint that the $(A[i] + i)$'s are different. The VARSELECT= option specifies the variable selection strategy to be first-in-first-out, the order in which the variables are encountered by the CLP procedure.

The following statements display the solution data set shown in Figure 3.2:

```
proc print data=out noobs label;
   label A1=a A2=b A3=c A4=d
         A5=e A6=f A7=g A8=h;
run;
```

Figure 3.2 A Solution to the Eight Queens Problem

a	b	c	d	e	f	g	h
1	5	8	6	3	7	2	4

The corresponding solution to the Eight Queens problem is displayed in Figure 3.3.

Figure 3.3 A Solution to the Eight Queens Problem

Syntax: CLP Procedure

The following statements are used in PROC CLP:

PROC CLP *options* ;
    ACTIVITY *activity specifications* ;
    ALLDIFF *alldiff constraints* ;
    ARRAY *array specifications* ;
    ELEMENT *element constraints* ;
    FOREACH *foreach constraints* ;
    GCC *global cardinality constraints* ;
    LINCON *linear constraints* ;
    REIFY *reify constraints* ;
    REQUIRES *resource requirement constraints* ;
    RESOURCE *resource specifications* ;
    SCHEDULE *schedule options* ;
    VARIABLE *variable specifications* ;

Functional Summary

The statements and options available with PROC CLP are summarized by purpose in Table 3.1.

Table 3.1 Functional Summary

Description	Statement	Option
Assignment Strategy Options		
Specifies the variable assignment strategy	PROC CLP	VARASSIGN=
Specifies the activity assignment strategy (Experimental)	SCHEDULE	ACTASSIGN=
Data Set Options		
Specifies the activity input data set (Experimental)	PROC CLP	ACTDATA=
Specifies the constraint input data set	PROC CLP	CONDATA=
Specifies the solution output data set	PROC CLP	OUT=
Specifies the resource input data set (Experimental)	PROC CLP	RESDATA=
Specifies the resource assignment data set (Experimental)	PROC CLP	SCHEDRES=
Specifies the time assignment data set (Experimental)	PROC CLP	SCHEDTIME=
Specifies the schedule output data set (Experimental)	PROC CLP	SCHEDULE=
General Options		
Specifies the upper bound on time (seconds)	PROC CLP	MAXTIME=
Suppresses preprocessing	PROC CLP	NOPREPROCESS
Permits preprocessing	PROC CLP	PREPROCESS
Specifies the units of MAXTIME	PROC CLP	TIMETYPE=
Implicitly defines Constraint data set variables	PROC CLP	USECONDATAVARS=

Table 3.1 *continued*

Description	Statement	Option
Output Control Options		
Finds all possible solutions	PROC CLP	FINDALLSOLNS
Specifies the number of solution attempts	PROC CLP	MAXSOLNS=
Indicates progress in log	PROC CLP	SHOWPROGRESS
Scheduling CSP-Related Statements (Experimental)		
Provides activity specifications	ACTIVITY	
Provides resource requirement specifications	REQUIRES	
Provides resource specifications	RESOURCE	
Provides scheduling parameters	SCHEDULE	
Scheduling: Resource Constraints (Experimental)		
Specifies the edge-finder consistency routines	SCHEDULE	EDGEFINDER=
Specifies the not-first edge-finder extension	SCHEDULE	NOTFIRST=
Specifies the not-last edge-finder extension	SCHEDULE	NOTLAST=
Scheduling: Temporal Constraints (Experimental)		
Specifies the activity duration	ACTIVITY	DURATION=
Specifies the activity finish lower bound	ACTIVITY	FGE=
Specifies the activity finish upper bound	ACTIVITY	FLE=
Specifies the activity start lower bound	ACTIVITY	SGE=
Specifies the activity start upper bound	ACTIVITY	SLE=
Specifies the schedule duration	SCHEDULE	DURATION=
Specifies the schedule finish	SCHEDULE	FINISH=
Specifies the schedule start	SCHEDULE	START=
Scheduling: Search Control Options (Experimental)		
Specifies the dead-end multiplier	PROC CLP	DM=
Specifies the number of allowable dead-ends per restart	PROC CLP	DPR=
Specifies the number of search restarts	PROC CLP	RESTARTS=
Selection Strategy Options		
Specifies the variable selection strategy	PROC CLP	VARSELECT=
Specifies the activity selection strategy (Experimental)	SCHEDULE	ACTSELECT=
Standard CSP Statements		
Specifies the all-different constraints	ALLDIFF	
Specifies the array specifications	ARRAY	
Specifies the element constraints	ELEMENT	
Specifies the for-each constraints	FOREACH	
Specifies the global cardinality constraints	GCC	
Specifies the linear constraints	LINCON	
Specifies the reified constraints	REIFY	
Specifies the variable specifications	VARIABLE	

PROC CLP Statement

> **PROC CLP** *options* ;

The following options can appear in the PROC CLP statement.

ACTDATA=_SAS-data-set_
ACTIVITY=_SAS-data-set_

>   identifies the input data set that defines the activities and temporal constraints. The temporal constraints consist of time-alignment-type constraints and precedence-type constraints. The format of the ACTDATA= data set is similar to that of the Activity data set used by the CPM procedure in SAS/OR software. The activities and time alignment constraints can also be specified directly by using the ACTIVITY statement without the need for a data set. The CLP procedure enables you to define activities by using a combination of the two specifications.

CONDATA=_SAS-data-set_

>   identifies the input data set that defines the constraints, variable types, and variable bounds. Currently, the CONDATA data set provides support for linear constraints only.

>   The linear constraints can also be specified in-line by using the LINCON statement. The CLP procedure enables you to define constraints by using a combination of the two specifications. When defining constraints, you must define the variables by using a VARIABLE statement or implicitly define them by specifying the USECONDATAVARS= option when using the CONDATA= data set. Note that variable bounds can be defined by using the VARIABLE statement, and any such definitions override those defined in the CONDATA= data set.

DM=_m_

>   specifies the dead-end multiplier for the scheduling CSP. The dead-end multiplier is used to determine the number of dead-ends that are permitted before triggering a complete restart of the search technique in a scheduling environment. The number of dead-ends is the product of the dead-end multiplier, m, and the number of unassigned activities. The default value is 0.15. This option is valid only with the SCHEDULE= option.

DOMAIN=_[lb, ub]_
DOM=_[lb, ub]_

>   specifies the global domain of all variables to be the closed interval [*lb*, *ub*]. You can override the global domain for a variable with a VARIABLE statement or the CONDATA= data set. The default domain is [0,∞].

DPR=_n_

>   specifies an upper bound on the number of dead-ends that are permitted before PROC CLP restarts or terminates the search, depending on whether or not a randomized search strategy is used. In the case of a nonrandomized strategy, *n* is an upper bound on the number of allowable dead-ends before terminating. In the case of a randomized strategy, *n* is an upper bound on the number of allowable dead-ends before restarting the search. The DPR= option has priority over the DM= option.

FINDALLSOLNS

ALLSOLNS

FINDALL

> attempts to find all possible solutions to the CSP. When a randomized search strategy is used, it is possible to rediscover the same solution and end up with multiple instances of the same solution. This is currently the case when you are solving a scheduling CSP. Therefore, this option is ignored when you are solving a scheduling CSP.

MAXSOLNS=*n*

> specifies the number of solution attempts to be generated for the CSP. The default value is 1. It is important to note, especially in the context of randomized strategies, that an attempt could result in no solution, given the current controls on the search mechanism, such as the number of restarts and the number of dead-ends permitted. As a result, the total number of solutions found might not match the MAXSOLNS= parameter.

MAXTIME=*n*

> specifies an upper bound on the number of seconds that are allocated for solving the problem. The type of time, either CPU time or real time, is determined by the value of the TIMETYPE= option. The default type is CPU time. The time specified by the MAXTIME= option is checked only once at the end of each iteration. Therefore, the actual running time can be longer than that specified by the MAXTIME= option. The difference depends on how long the last iteration takes. The default value of MAXTIME= is ∞. If you do not specify this option the procedure does not stop based on the amount of time elapsed.

NOPREPROCESS

> suppresses any preprocessing that would typically be performed for the problem.

OUT=*SAS-data-set*

> identifies the output data set that contains one or more solutions to a standard CSP, if one exists. Each observation in the OUT= data set corresponds to a solution of the CSP. The number of solutions that are generated can be controlled using the MAXSOLNS= option in the PROC CLP statement.

PREPROCESS

> permits any preprocessing that would typically be performed for the problem.

RESDATA=*SAS-data-set*

RESIN=*SAS-data-set*

> identifies the input data set that defines the resources and their attributes such as capacity and resource pool membership. This information can be used in lieu of, or in combination with, the RESOURCE statement.

RESTARTS=*n*

> specifies the number of restarts of the randomized search technique before terminating the procedure. The default value is 3.

SCHEDRES=*SAS-data-set* ***(Experimental)***

> identifies the output data set that contains the solutions to scheduling CSPs. This data set contains the resource assignments of activities.

SCHEDTIME=*SAS-data-set* *(Experimental)*

  identifies the output data set that contains the solutions to scheduling CSPs. This data set contains the time assignments of activities.

SCHEDULE=*SAS-data-set* *(Experimental)*

SCHEDOUT=*SAS-data-set* *(Experimental)*

  identifies the output data set that contains the solutions to a scheduling CSP, if any exist. This data set contains both the time and resource assignment information. There are two types of observations identified by the value of the OBSTYPE variable. Observation with OBSTYPE= "TIME" corresponds to time assignment, and observation with OBSTYPE= "RESOURCE" corresponds to resource assignment. The maximum number of solutions can be controlled by using the MAXSOLNS= option in the PROC CLP statement.

SHOWPROGRESS

  prints a message to the log whenever a solution has been found. When a randomized strategy is used, the number of restarts and dead-ends that were required are also printed to the log.

TIMETYPE=CPU | REAL

  specifies the units for the value of the MAXTIME= option. The value of this option determines whether such time is CPU time or real time. The default value of this option is CPU.

USECONDATAVARS=0 | 1

  specifies whether the numeric variables in the CONDATA= data set, with the exception of any reserved variables, are implicitly defined or not. A value of 1 indicates they are implicitly defined, in which case a VARIABLE statement is not necessary to define the variables in the data set. The default value is 0. Currently, _RHS_ is the only reserved numeric variable.

VARASSIGN=*keyword*

  specifies the value selection strategy. Currently, there is only one value selection strategy. The MIN strategy selects the minimum value from the domain of the selected variable. To assign activities, use the ACTASSIGN= option in the SCHEDULE statement.

VARSELECT=*keyword*

  specifies the variable selection strategy. Both static and dynamic strategies are available.

  Static strategies are as follows:

   • FIFO, which uses the first-in-first-out ordering of the variables as encountered by the procedure

   • MAXCS, which selects the variable with the maximum number of constraints

  Dynamic strategies are as follows:

   • MINR, which selects the variable with the smallest range (that is, the minimum value of upper bound minus lower bound)

   • MAXC, which selects the variable with the largest number of active constraints

   • MINRMAXC, which selects the variable with the smallest range, breaking ties by selecting one with the largest number of active constraints

The dynamic strategies embody the "fail first principle" (FFP) of Haralick and Elliot (1980), which suggests that "To succeed, try first where you are most likely to fail." The default activity selection strategy is MINR. To set the strategy for selecting activities, use the ACTSELECT= option in the SCHEDULE statement.

ACTIVITY Statement

> **ACTIVITY** *specification-1 < . . . specification-n >* **;**

An ACTIVITY *specification* can be one of the following types:

> *activity < = (<* **DUR=** *> duration < altype=aldate . . . >) >*

> *(activity_list) < = (<* **DUR=** *> duration < altype=aldate . . . >) >*

where *duration* is the activity duration and *altype* is a keyword that specifies an alignment-type constraint on the activity (or activities) with respect to the value given by *aldate*.

The ACTIVITY statement defines one or more activities and the attributes of each activity, such as the duration and any temporal constraints of the time-alignment-type. The activity duration can take nonnegative integer values. The default duration is 0.

Valid *altype* keywords are as follows:

- **SGE**, start greater than or equal to *aldate*

- **SLE**, start less than or equal to *aldate*

- **FGE**, finish greater than or equal to *aldate*

- **FLE**, finish less than or equal to *aldate*

You can specify any combination of the preceding keywords. For example, to define activities A1, A2, A3, B1, and B3 with duration 3, and to set the start time of these activities equal to 10, specify the following:

```
activity (A1-A3 B1 B3) = ( dur=3 sge=10 sle=10 );
```

If an activity appears in more than one ACTIVITY statement, only the first activity definition is honored. Additional specifications are ignored.

You can alternatively use the ACTDATA= data set to define activities, durations, and temporal constraints. In fact, you can specify both an ACTIVITY statement and an ACTDATA= data set. You must use an ACTDATA= data set to define precedence-related temporal constraints. One of SCHEDULE=, SCHEDRES=, or SCHEDTIME= must be specified when the ACTIVITY statement is used.

ALLDIFF Statement

ALLDIFF *(variable_list-1)* < . . . *(variable_list-n)* > **;**

ALLDIFFERENT *(variable_list-1)* < . . . *(variable_list-n)* > **;**

The ALLDIFF statement can have multiple specifications. Each specification defines a unique global constraint on a set of variables, requiring all of them to be different from each other. A global constraint is equivalent to a conjunction of elementary constraints.

For example, the statements

```
var (X1-X3) A B;
alldiff (X1-X3) (A B);
```

are equivalent to

$$
\begin{aligned}
X1 &\neq X2 \text{ AND} \\
X2 &\neq X3 \text{ AND} \\
X1 &\neq X3 \text{ AND} \\
A &\neq B
\end{aligned}
$$

If the variable list is empty, the ALLDIFF constraint applies to all the variables declared in any VARIABLE statement.

ARRAY Statement

ARRAY *specification-1* < . . . *specification-n* > **;**

An ARRAY *specification* is in a form as follows:

name[dimension](variables)

The ARRAY statement is used to associate a *name* with a list of *variables*. Each of the *variables* in the variable list must be defined using a VARIABLE statement or implicitly defined using the CONDATA= data set. The ARRAY statement is required when you are specifying a constraint by using the FOREACH statement.

ELEMENT Statement

ELEMENT *element_constraint-1* < . . . *element_constraint-n* > **;**

An *element_constraint* is specified in the following form:

(index variable, (integer list), variable)

The ELEMENT statement specifies one or more element constraints. An element constraint enables you to define dependencies, not necessarily functional, between variables. The statement

ELEMENT$(I, (L), V)$

sets the variable V to be equal to the Ith element in the list L. The list of integers $L = (v_1, ..., v_n)$ is a list of values that the variable V can take and are not necessarily distinct. The variable I is the index variable and its domain is considered to be $[1, n]$. Each time the domain of I is modified, the domain of V is updated and vice versa.

An element constraint enforces the following propagation rules:

$$V = v \Leftrightarrow I \in \{i_1, ..., i_m\}$$

where v is a value in the list L and $i_1, ..., i_m$ are all the indices in L whose value is v.

The following statements use the element constraint to implement the quadratic function $y = x^2$:

```
proc clp out=clpout;
   var x=[1,5] y=[1,25];
   element (x,(1,  4,  9, 16,  25), y);
run;
```

An element constraint is equivalent to a conjunction of reify and linear constraints. For example, the preceding statements are equivalent to:

```
proc clp out=clpout;
        var x=[1,5] y=[1,25] (R1-R5)=[0,1];
        reify R1: (x=1);
        reify R1: (y=1);
        reify R2: (x=2);
        reify R2: (y=4);
        reify R3: (x=3);
        reify R3: (y=9);
        reify R4: (x=4);
        reify R4: (y=16);
        reify R5: (x=5);
        reify R5: (y=25);
        lincon R1 + R2 + R3 + R4 + R5 = 1;
   run;
```

Element constraints can also be used to define positional mappings between two variables. For example, suppose the function $y = x^2$ is defined on only odd numbers in the interval $[-5, 5]$. You can model this by using two element constraints and an artificial index variable:

```
element (i, ( -5, -3, -1, 1, 3,  5), x)
        (i, ( 25,  9,  1, 1, 9, 25), y);
```

The list of values L can also be specified by using a convenient syntax of the form *start* TO *end* or *start* TO *end* BY *increment*. For example, the previous element specification is equivalent to:

```
element (i, ( -5 to 5 by 2), x)
        (i, ( 25,  9,  1, 1, 9, 25), y);
```

FOREACH Statement

> **FOREACH** *(array, type, < offset >)* **;**

where *array* must be defined by using an ARRAY statement, *type* is a keyword that determines the type of the constraint, and *offset* is an integer.

The FOREACH statement iteratively applies a constraint over an array of variables. The type of the constraint is determined by *type*. Currently, the only valid *type* keyword is DIFF. The optional *offset* parameter is an integer and is interpreted in the context of the constraint type. The default value of *offset* is zero.

The FOREACH statement that corresponds to the DIFF keyword iteratively applies the following constraint to each pair of variables in the array:

$$variable\_i + offset \times i \ \neq \ variable\_j + offset \times j \quad \forall \ i \neq j, \ i, j = 1, \ldots, array\_dimension$$

For example, the constraint that all $(A[i] - i)$'s are pairwise distinct for an array A is expressed as

```
foreach (A, diff, -1);
```

GCC Statement

> **GCC** *global_cardinality_constraint-1 < ... global_cardinality_constraint-n >* **;**

where *global_cardinality_constraint* is specified in the following form:

> *(variables)* = ($(v_1, l_1, u_1) < \ldots (v_n, l_n, u_n) >$ < **DL=** *dl* > < **DU=** *du* >)

v_i is a value in the domain of one of the variables, and l_i and u_i are the lower and upper bounds on the number of variables assigned to v_i. The values of *dl* and *du* are the lower and upper bounds on the number of variables assigned to values in D outside of $\{v_1, \ldots, v_n\}$.

The GCC statement specifies one or more global cardinality constraints. A *global cardinality constraint* (GCC) is a constraint that consists of a set of variables $\{x_1, \ldots, x_n\}$ and for each value v in $D = \bigcup_{i=1,\ldots,n} Dom(x_i)$, a pair of numbers l_v and u_v. A GCC is satisfied if and only if the number of times that a value v in D is assigned to the variables $x_1, \ldots, x_n$ is at least l_v and at most u_v.

For example, the constraint that is specified with the statements

```
var (x1-x6) = [1, 4];
gcc(x1-x6) = ((1, 1, 2) (2, 1, 3) (3, 1, 3) (4, 2, 3));
```

expresses that at least one but no more than two variables in $x_1, \ldots, x_6$ can have value 1, at least one and no more than three variables can have value 2 (or value 3), and at least two and no more than three variables can have value 4. For example, an assignment $x_1 = 1, x_2 = 1, x_3 = 2, x_4 = 3, x_5 = 4$, and $x_6 = 4$ satisfies the constraint.

If a global cardinality constraint has common lower or upper bounds for many of the values in D, the DL= and DU= options can be used to specify the common lower and upper bounds.

For example, the previous specification could also be written as

```
gcc(x1-x6) = ((1, 1, 2) (4, 2, 3) DL=1 DU=3);
```

You can also specify missing values for the lower and upper bounds. The values of *dl* and *du* are substituted as appropriate. The previous example can also be expressed as

```
gcc(x1-x6) = ((1, ., 2) (4, 2, .) DL=1 DU=3);
```

The following statements specify that each of the values in $\{1, \ldots, 9\}$ can be assigned to at most one of the variables $x_1, \ldots, x_9$:

```
var (x1-x9) = [0, 9];
gcc(x1-x9) = (DL=0 DU=1);
```

Note that the preceding global cardinality constraint is equivalent to the all-different constraint that is expressed as:

```
var (x1-x9) = [0, 9];
alldiff(x1-x9);
```

If you do not specify the DL= and DU= options, the default lower and upper bound for any value in D that does not appear in the (v, l, u) format is 0 and the number of variables in the constraint, respectively.

The global cardinality constraint also provides a convenient way of defining disjoint domains for a set of variables. For example, the following syntax limits assignment of the variables $x_1, \ldots, x_9$ to even numbers between 0 and 10:

```
var (x1-x9) = [0, 10];
gcc(x1-x9) = ((1, 0, 0) (3, 0, 0) (5, 0, 0) (7, 0, 0) (9, 0, 0));
```

If the variable list is empty, the GCC constraint applies to all the variables declared in any VARIABLE statement.

LINCON Statement

> **LINCON** *linear_constraint-1* < . . . ,*linear_constraint-n* > **;**

> **LINEAR** *linear_constraint-1* < . . . ,*linear_constraint-n* > **;**

where *linear_constraint* has the form

> *linear_expression-l type linear_expression-r*

where *linear_expression* has the form

> < **+**|**-** >*linear_term-1* < . . . , **(+|-)** *linear_term-n* >

where *linear_term* has the form

> **(***variable* | *number*< ***** *variable* >**)**

The keyword *type* can be one of the following:

<, <=, =, >=, >, <>, LT, LE, EQ, GE, GT, NE

The LINCON statement allows for a very general specification of linear constraints. In particular, it allows for specification of the following types of equality or inequality constraints:

$$\sum_{j=1}^{n} a_{ij} x_j \;\{\leq \mid < \mid = \mid \geq \mid > \mid \neq\}\; b_i \quad \text{for } i = 1, \ldots, m$$

For example, the constraint $4x_1 - 3x_2 = 5$ can be expressed as

```
var x1 x2;
lincon 4 * x1 - 3 * x2 = 5;
```

and the constraints

$$
\begin{aligned}
10x_1 - x_2 &\geq 10 \\
x_1 + 5x_2 &\neq 15
\end{aligned}
$$

can be expressed as

```
var x1 x2;
lincon 10 <= 10 * x1 - x2,
       x1 + 5 * x2 <> 15;
```

Note that variables can be specified on either side of an equality or inequality in a LINCON statement. Linear constraints can also be specified by using the CONDATA= data set.

Regardless of the specification, you must define the variables by using a VARIABLE statement or implicitly by specifying the USECONDATAVARS= option.

Note that user-specified scalar values are subject to rounding based upon a platform-dependent tolerance.

REIFY Statement

REIFY *reify_constraint-1 < . . . reify_constraint-n >* ;

where *reify_constraint* is specified in the following form:

variable : constraint

The REIFY statement associates a binary variable with a constraint. The value of the binary variable is 1 or 0 depending on whether the constraint is satisfied or not, respectively. The constraint is said to be *reified*, and the binary variable is referred to as the *control variable*. Currently, the only type of constraint that can be reified is the linear constraint, which should have the same form as *linear_constraint* defined in the LINCON statement. As with the other variables, the control variable must also be defined in a VARIABLE statement or in the CONDATA= data set.

The REIFY statement provides a convenient mechanism for expressing logical constraints, such as disjunctive and implicative constraints. For example, the disjunctive constraint

$$(3x + 4y < 20) \vee (5x - 2y > 50)$$

can be expressed with the following statements:

```
var x y p q;
reify p: (3 * x + 4 * y < 20) q: (5 * x - 2 * y > 50);
lincon p + q >= 1;
```

The binary variables p and q reify the linear constraints

$$3x + 4y < 20$$

and

$$5x - 2y > 50$$

respectively. The following linear constraint enforces the desired disjunction:

$$p + q \geq 1$$

The implication constraint

$$(3x + 4y < 20) \Rightarrow (5x - 2y > 50)$$

can be enforced with the linear constraint

$$q - p \geq 0$$

The REIFY constraint can also be used to express a constraint that involves the absolute value of a variable. For example, the constraint

$$|X| = 5$$

can be expressed with the following statements:

```
var x p q;
reify p: (x = 5) q: (x = -5);
lincon p + q = 1;
```

REQUIRES Statement

> **REQUIRES** *resource_constraint-1 <...resource_constraint-n> ;*

where *resource_constraint* is specified in the following form:

> *activity_specification = (resource_specification) < qty = q >*

where

> *activity_specification: (activity | activity-1 <...activity-m>)*

and

> *resource_specification: (resource-1 < qty = r_1 > <...(, | **OR**) resource-l < qty=r_l >>)*

activity_specification is a single activity or a list of activities that requires *q* units of the resource identified in *resource_specification*. *resource_specification* is a single resource or a list of resources, representing a choice of resource, along with the equivalent required quantities for each resource. The default value of r_i is 1. Alternate resource requirements are separated by a comma (,) or the keyword OR. The *qty=* parameter outside the *resource_specification* acts as a multiplier to the *qty=* parameters inside the *resource_specification*.

The REQUIRES statement defines the potential activity assignments with respect to the pool of resources. If an activity is not defined, the REQUIRES statement implicitly defines the activity.

You can also define resource constraints by using the Activity and Resource data sets in lieu of, or in conjunction with, the REQUIRES statement. Any resource constraints that are defined for an activity by using a REQUIRES statement override all resource constraints for that activity that are defined by using the Activity and Resource data sets.

The following statements illustrate how you would use a REQUIRES statement to specify that activity A requires resource R:

```
activity A;
resource R;
requires A = (R);
```

In order to specify that activity A requires two units of the resource R, you would add the *qty=* keyword as in the following example:

```
requires A = (R qty=2);
```

In certain situations, the assignment might not be established in advance and there might be a set of possible alternates that can satisfy the requirements of an activity. This scenario can be defined by using multiple *resource-specifications* separated by commas or the keyword OR. For example, if the activity A needs either two units of the resource R1 or one unit of the resource R2, you could use the following statement:

```
requires A = (R1 qty=2, R2);
```

The equivalent statement using the keyword OR is

```
requires A = (R1 qty=2 or R2);
```

It is important to note that resources specified in a single resource constraint are disjunctive and not conjunctive. The activity is satisfied by exactly one of the resources rather than a combination of resources. For example, the following statement specifies that the possible resource assignment for activity A is either four units of R1 or two units of R2:

```
requires A = (R1 qty=2 or R2) qty=2;
```

The preceding statement does not, for example, result in an assignment of two units of the resource R1 and one unit of R2.

In order to model conjunctive resources by using a REQUIRES statement, such as when an activity requires more than one resource simultaneously, you need to define multiple resource constraints. For example, if activity A requires both resource R1 and resource R2, you can model it as follows:

```
requires A = (R1)  A = (R2);
```

or

```
requires A = (R1);
requires A = (R2);
```

If multiple activities have the same resource requirements, you can use an activity list for specifying the constraints instead of having separate constraints for each activity. For example, if activities A and B require resource R1 or resource R2, the specification

```
requires (A B) = (R1, R2);
```

is equivalent to

```
requires A = (R1, R2);
requires B = (R1, R2);
```

RESOURCE Statement

> **RESOURCE** *resource_specification-1* < ... *resource_specification-n* > ;

where *resource_specification* is specified in the following form:

> *resource* | *(resource-1* < ... *resource-m* >) < =*(capacity)* >

The RESOURCE statement specifies the names and capacities of all resources that are available to be assigned to any defined activities. For example, the following statement specifies that there are two units of the resource R1 and one unit of the resource R2.

```
resource R1=(2) R2;
```

The capacity of a resource can take nonnegative integer values. The default capacity is 1, which corresponds to a unary resource.

SCHEDULE Statement

>**SCHEDULE** *options* ;

>**SCHED** *options* ;

The following options can appear in the SCHEDULE statement.

ACTASSIGN=*keyword*

> specifies the activity assignment strategy subject to the activity selection strategy, which is specified by the ACTSELECT= option. After an activity has been selected, the activity assignment strategy determines a start time and a set of resources (empty if the activity has no resource requirements) for the selected activity. The interpretation of the assignment strategy depends on whether the activity selection strategy has been specified as RJRAND or not. The activity is assigned its earliest possible start time unless ACTSELECT=RJRAND; otherwise, the activity is assigned its latest possible start time.

> Figure 3.4 illustrates possible start times for a single activity, which requires one of the resources R1, R2, R3, R4, R5, or R6. The bars depict the possible start times that are supported by each of the resources for the duration of the activity.

Figure 3.4 Potential Activity Start Times
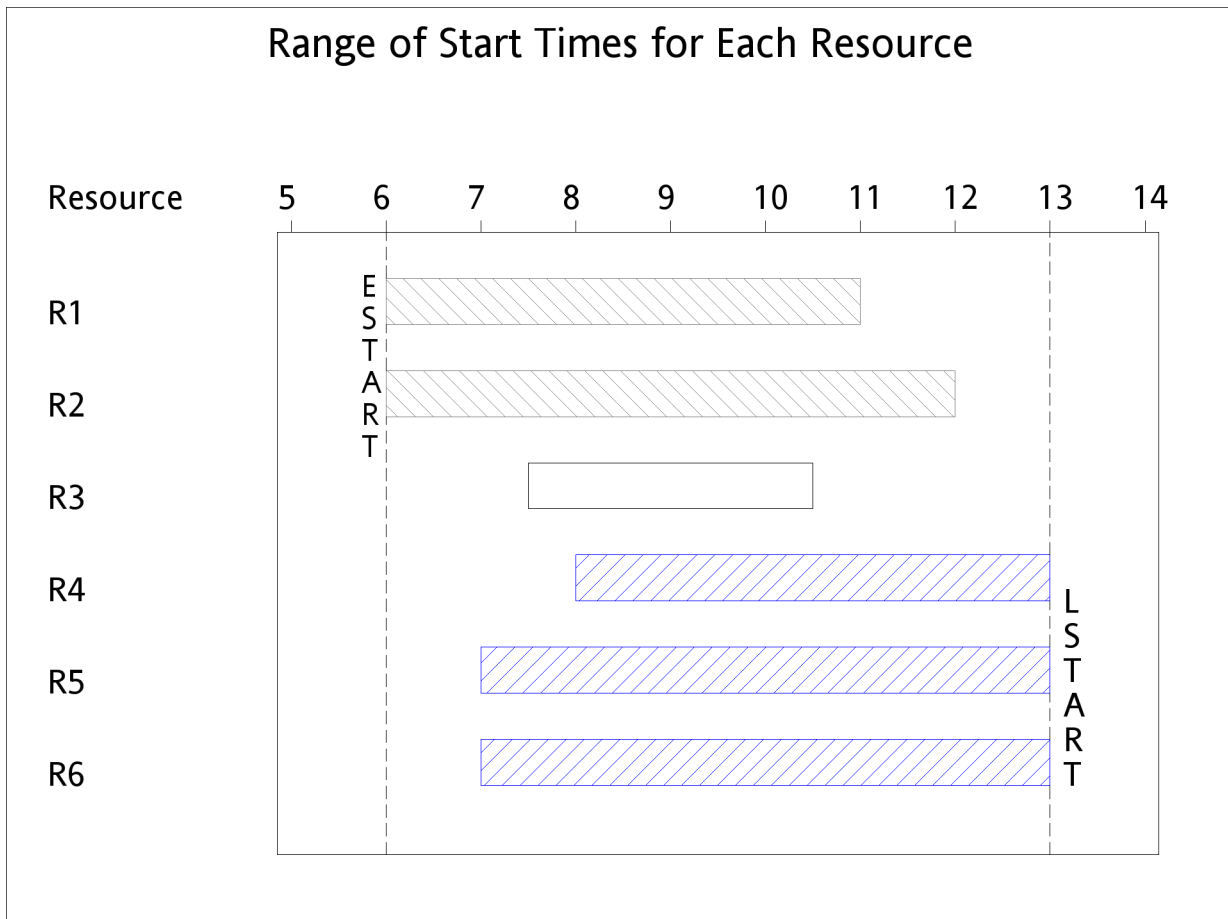
For example, if ACTSELECT=LJRAND, the activity is assigned a start time of 6 and one of R1 or R2 is assigned. On the other hand, if ACTSELECT=RJRAND, the activity is assigned a start time of 13 and one of R4, R5, or R6 is assigned.

If the activity has any resource requirements, then the activity is assigned a set of resources as follows:

RAND   randomly selects a set of resources that support the selected start time for the activity.

                         In Figure 3.4, if the activity start time is set to 6, the strategy randomly selects between R1 and R2. Otherwise, the strategy randomly selects among R4, R5, and R6.

MAXTW | MAXLS   selects the set of resources that supports the assigned start time and affords the maximum time window of availability for the activity. Ties are broken randomly.

                         In Figure 3.4, if the activity start time is set to 6, the resources that support the selected start time are R1 and R2. Since R1 has a smaller time window, the strategy selects R2. On the other hand, if the activity start time is set to 13, the resources that support the selected start time are R4, R5, and R6. Because R4 has a smaller time window than R5 or R6, the strategy randomly selects between R5 and R6.

The default activity assignment strategy is RAND. For assigning variables, use the VARASSIGN= option in the PROC CLP statement.

ACTSELECT=_keyword_

specifies the activity selection strategy. The activity selection strategy can be randomized or deterministic.

The following selection strategies use a random heuristic to break ties:

LJRAND | RAND   selects an activity at random from those that begin prior to the earliest early finish time. This strategy was proposed by Nuijten (1994).

MAXD   selects an activity at random from those that begin prior to the earliest early finish time and that have maximum duration.

MINA   selects an activity at random from those that begin prior to the earliest early finish time and that have the minimum number of resource assignments.

MINLS   selects an activity at random from those that begin prior to the earliest early finish time and that have a minimum late start time.

RJRAND   selects an activity at random from those that finish after the latest late start time.

The following are deterministic selection strategies:

DET   selects the first activity that begins prior to the earliest activity finish time.

DMINLS   selects the activity with the earliest late start time.

The first activity is defined according to its appearance in the following order of precedence:

1. ACTIVITY statement
2. REQUIRES statement
3. ACTDATA= data set

The default activity selection strategy is RAND. For selecting variables, use the VARSELECT= option in the PROC CLP statement.

DURATION=*dur*

SCHEDDUR=*dur*

DUR=*dur*

specifies the duration of the schedule. The DURATION= option imposes a constraint that the duration of the schedule does not exceed the specified value.

EDGEFINDER *<=eftype>*

EDGE *<=eftype>*

activates the edge-finder consistency routines for scheduling CSPs. By default, the EDGEFINDER= option is inactive. Specifying the EDGEFINDER= option determines whether an activity must be the first or the last to be processed from a set of activities that require a given resource or set of resources and prunes the domain of the activity appropriately.

Valid values for the *eftype* keyword are FIRST, LAST, or BOTH. Note that *eftype* is an optional argument, and that specifying EDGEFINDER by itself is equivalent to specifying EDGEFINDER=LAST. The interpretation of each of these keywords is described as follows:

- FIRST: The edge-finder algorithm attempts to determine whether an activity must be processed first from a set of activities that require a given resource or set of resources and prunes its domain appropriately.
- LAST: The edge-finder algorithm attempts to determine whether an activity must be processed last from a set of activities that require a given resource or set of resources and prunes its domain appropriately.
- BOTH: This is equivalent to specifying both FIRST and LAST. The edge-finder algorithm attempts to determine which activities must be first and which activities must be last, and updates their domains as necessary.

There are several extensions to the edge-finder consistency routines. These extensions are invoked by using the NOTFIRST= and NOTLAST= options in the SCHEDULE statement. For more information about options that are related to edge-finder consistency routines, see the section "Edge Finding" on page 50.

FINISH=_finish_

END=_finish_

FINISHBEFORE=_finish_

>   specifies the finish time for the schedule. The schedule finish time is an upper bound on the finish time of each activity (subject to time, precedence, and resource constraints). If you want to impose a tighter upper bound for an activity, you can do so either by using the FLE= specification in an ACTIVITY statement or by using the _ALIGNDATE_ and _ALIGNTYPE_ variables in the ACTDATA= data set.

NOTFIRST=_level_

NF=_level_

>   activates an extension of the edge-finder consistency routines for scheduling CSPs. By default, the NOTFIRST= option is inactive. Specifying the NOTFIRST= option determines whether an activity cannot be the first to be processed from a set of activities that require a given resource or set of resources and prunes its domain appropriately.

>   The argument _level_ is numeric and indicates the level of propagation. Valid values are 1, 2, or 3, with a higher number reflecting more propagation. More propagation usually comes with a higher performance cost; the challenge is to strike the right balance. Specifying the NOTFIRST= option implicitly turns on the EDGEFINDER=LAST option because the latter is a special case of the former.

>   The corresponding NOTLAST= option determines whether an activity cannot be the last to be processed from a set of activities that require a given resource or set of resources.

>   For more information about options that are related to edge-finder consistency routines, see the section "Edge Finding" on page 50.

NOTLAST=_level_

NL=_level_

>   activates an extension of the edge-finder consistency routines for scheduling CSPs. By default, the NOTLAST= option is inactive. Specifying the NOTLAST= option determines whether an activity cannot be the last to be processed from a set of activities that require a given resource or set of resources and prunes its domain appropriately.

>   The argument _level_ is numeric and indicates the level of propagation. Valid values are 1, 2, or 3, with a higher number reflecting more propagation. More propagation usually comes with a higher performance cost; the challenge is to strike the right balance. Specifying the NOTLAST= option implicitly turns on the EDGEFINDER=FIRST option because the latter is a special case of the former.

>   The corresponding NOTFIRST= option determines whether an activity cannot be the first to be processed from a set of activities requiring a given resource or set of resources.

>   For more information about options that are related to edge-finder consistency routines, see the section "Edge Finding" on page 50.

START=*start*

BEGIN=*start*

STARTAFTER=*start*

> specifies the start time for the schedule. The schedule start time is a lower bound on the start time of each activity (subject to time, precedence, and resource constraints). If you want to impose a tighter lower bound for an activity, you can do so either by using the SGE= specification in an ACTIVITY statement or by using the _ALIGNDATE_ and _ALIGNTYPE_ variables in the ACTDATA= data set.

VARIABLE Statement

> **VARIABLE** *var_specification-1* < *... var_specification-n* > **;**

> **VAR** *var_specification-1* < *... var_specification-n* > **;**

A *var_specification* can be one of the following types:

> *variable* < *=[lower-bound* < *, upper-bound* >*]* >

> *(variables)* < *=[lower-bound* < *, upper-bound* >*]* >

The VARIABLE statement declares all variables that are to be considered in the CSP and, optionally, defines their domains. Any variable domains defined in a VARIABLE statement override the global variable domains that are defined by using the DOMAIN= option in the PROC CLP statement in addition to any bounds that are defined by using the CONDATA= data set. If *lower-bound* is specified and *upper-bound* is omitted, the corresponding variables are considered as being assigned to *lower-bound*. The values of *lower-bound* and *upper-bound* can also be specified as missing, in which case the appropriate values from the DOMAIN= specification are substituted.

Details: CLP Procedure

Modes of Operation

The CLP procedure can be invoked in either of the following modes:

- The *standard mode* gives you access to all-different constraints, element constraints, GCC constraints, linear constraints, reify constraints, ARRAY statements, and FOREACH statements. In standard mode, the decision variables are one-dimensional; a variable is assigned an integer in a solution.

- The *scheduling mode* gives you access to more scheduling-specific constraints, such as temporal constraints (precedence and time) and resource constraints. In scheduling mode, the

variables are typically multidimensional; a variable is assigned a start time and possibly a set of resources in a solution. In scheduling mode, the variables are referred to as activities, and the solution is referred to as a schedule. Scheduling mode is experimental in the current version of the CLP procedure.

Selecting the Mode of Operation

The CLP procedure requires the specification of an output data set to store one or more solutions to the CSP. There are four possible output data sets: the Solution data set (specified using the OUT= option in the PROC CLP statement), which corresponds to the standard mode of operation, and one or more Schedule data sets (specified using the SCHEDULE=, SCHEDRES=, or SCHEDTIME= options in the PROC CLP statement), which correspond to the scheduling mode of operation. The mode is determined by which output data set has been specified. If an output data set is not specified, the procedure terminates with an error message. If both types of output data sets have been specified, the schedule-related data sets are ignored.

Constraint Data Set

The Constraint data set defines linear constraints, variable types, and bounds on variable domains. You can use a Constraint data set in lieu of, or in combination with, a LINCON or a VARIABLE statement (or both) in order to define linear constraints, variable types, and variable bounds. The Constraint data set is specified by using the CONDATA= option in the PROC CLP statement.

The Constraint data set must be in dense input format. In this format, a model's columns appear as variables in the input data set and the data set must contain the _TYPE_ variable, at least one numeric variable, and any reserved variables. Currently, the only reserved variable is the _RHS_ variable. If this requirement is not met, the CLP procedure terminates. The _TYPE_ variable is a character variable that tells the CLP procedure how to interpret each observation. The CLP procedure recognizes the following keywords as valid values for the _TYPE_ variable: EQ, LE, GE, NE, LT, GT, LOWERBD, UPPERBD, BINARY, and FIXED. An optional character variable, _ID_, can be used to name each row in the Constraint data set.

Linear Constraints

For the _TYPE_ values EQ, LE, GE, NE, LT, and GT, the corresponding observation is interpreted as a linear constraint. The _RHS_ variable is a numeric variable that contains the right-hand-side coefficient of the linear constraint. Any numeric variable other than _RHS_ that appears in a VARIABLE statement is interpreted as a structural variable for the linear constraint.

The _TYPE_ values are defined as follows:

EQ (=)   defines a linear equality of the form

$$\sum_{j=1}^{n} a_{ij}x_j = b_i$$

LE (<=)   defines a linear inequality of the form

$$\sum_{j=1}^{n} a_{ij}x_j \leq b_i$$

GE (>=)   defines a linear inequality of the form

$$\sum_{j=1}^{n} a_{ij}x_j \geq b_i$$

NE (<>)   defines a linear disequation of the form

$$\sum_{j=1}^{n} a_{ij}x_j \neq b_i$$

LT (<)   defines a linear inequality of the form

$$\sum_{j=1}^{n} a_{ij}x_j < b_i$$

GT (>)   defines a linear inequality of the form

$$\sum_{j=1}^{n} a_{ij}x_j > b_i$$

Domain Bounds

The keywords LOWERBD and UPPERBD specify additional lower bounds and upper bounds, respectively, on the variable domains. In an observation where the _TYPE_ variable is equal to LOWERBD, a nonmissing value for a decision variable is considered to be a lower bound for that variable. Similarly, in an observation where the _TYPE_ variable is equal to UPPERBD, a nonmissing value for a decision variable is considered to be an upper bound for that variable. Note that lower and upper bounds defined in the Constraint data set are overridden by lower and upper bounds that are defined by using a VARIABLE statement.

Variable Types

The keywords BINARY and FIXED specify numeric variable types. If the value of _TYPE_ is BINARY for an observation, then any decision variable with a nonmissing entry for the observation is interpreted as being a binary variable with domain {0,1}. If the value of _TYPE_ is FIXED for an observation, then any decision variable with a nonmissing entry for the observation is interpreted as being assigned to that nonmissing value. In other words, if the value of the variable X is c in an observation for which _TYPE_ is FIXED, then the domain of X is considered to be the singleton $\{c\}$. The value c should belong to the domain of X, or the problem is deemed infeasible.

Variables in the CONDATA= Data Set

Table 3.2 lists all the variables that are associated with the Constraint data set and their interpretations by the CLP procedure. For each variable, the table also lists its type (C for character, N for numeric), the possible values it can assume, and its default value.

Table 3.2   Constraint Data Set Variables

Name	Type	Description	Allowed Values	Default
TYPE	C	Observation type	EQ, LE, GE, NE, LT, GT, LOWERBD, UPPERBD, BINARY, FIXED	
RHS	N	Right-hand-side coefficient		0
_ID_	C	Observation name (optional)		
Any numeric variable other than _RHS_	N	Structural variable		

Solution Data Set

In order to solve a standard (nonscheduling) type CSP, you need to specify a Solution data set by using the OUT= option in the PROC CLP statement. The Solution data set contains all the solutions that have been determined by the CLP procedure. You can specify an upper bound on the number of solutions by using the MAXSOLNS= option in the PROC CLP statement. If you prefer that PROC CLP determine all possible solutions instead, you can specify the FINDALLSOLNS option in the PROC CLP statement.

The Solution data set contains as many decision variables as have been defined in the CLP procedure invocation. Every observation in the Solution data set corresponds to a solution to the CSP. If a Constraint data set has been specified, then any variable formats and variable labels from the Constraint data set carry over to the Solution data set.

Activity Data Set

You can use an Activity data set in lieu of, or in combination with, an ACTIVITY statement to define activities and constraints that relate to the activities. The Activity data set is similar to the Activity data set of the CPM procedure in SAS/OR software and is specified by using the ACTDATA= option in the PROC CLP statement.

The Activity data set enables you to define an activity, its domain, temporal constraints, and resource constraints. The temporal constraints can be either time-alignment-type or precedence-type constraints. The Activity data set requires at least two variables: one to determine the activity, and another to determine its duration. The procedure terminates if it cannot find the required variables. The activity is determined with the _ACTIVITY_ variable, which must be character, and the duration is determined with the _DURATION_ variable, which must be numeric. You can define temporal constraints and resource constraints by including additional variables.

Time Alignment Constraints

The _ALIGNDATE_ and _ALIGNTYPE_ variables enable you to define time-alignment-type constraints. The _ALIGNTYPE_ variable defines the type of the alignment constraint for the activity that is named in the _ACTIVITY_ variable with respect to the _ALIGNDATE_ variable. If the _ALIGNDATE_ variable is not present in the Activity data set, the _ALIGNTYPE_ variable is ignored. Similarly, _ALIGNDATE_ is ignored when _ALIGNTYPE_ is not present. The _ALIGNDATE_ variable can take nonnegative integer values. The _ALIGNTYPE_ variable can take the values shown in Table 3.3.

Table 3.3 Valid Values for the _ALIGNTYPE_ Variable

Value	Type of Alignment
SEQ	Start equal to
SGE	Start greater than or equal to
SLE	Start less than or equal to
FEQ	Finish equal to
FGE	Finish greater than or equal to
FLE	Finish less than or equal to

Precedence Constraints

The _SUCCESSOR_ variable enables you to define precedence-type relationships between activities by using AON (activity-on-node) format. The _SUCCESSOR_ variable is a character variable. The _LAG_ variable defines the lag type of the relationship. By default, all precedence relationships are considered to be *finish-to-start* (*FS*). An FS type of precedence relationship is also referred to as a *standard* precedence constraint. All other types of precedence relationships are considered to be nonstandard precedence constraints. The _LAGDUR_ variable specifies the lag duration. By default, the lag duration is zero.

For each (activity, successor) pair, you can define a lag type and a lag duration. Consider a pair of activities (A, B) with a lag duration given by *lagdur*. The interpretation of each of the different lag types is given in Table 3.4.

Table 3.4 Valid Values for the _LAG_ Variable

Lag Type	Interpretation
FS	Finish A + lagdur $\leq$ Start B
SS	Start A + lagdur $\leq$ Start B
FF	Finish A + lagdur $\leq$ Finish B
SF	Start A + lagdur $\leq$ Finish B
FSE	Finish A + lagdur = Start B
SSE	Start A + lagdur = Start B
FFE	Finish A + lagdur = Finish B
SFE	Start A + lagdur = Finish B

The first four lag types (FS, SS, FF, and SF) are also referred to as *finish-to-start*, *start-to-start*, *finish-to-finish*, and *start-to-finish*, respectively. The next four types (FSE, SSE, FFE, and SFE) are stricter versions of FS, SS, FF, and SF, respectively. The first four types impose a lower bound on the start and finish times of B, while the last four types force the start and finish times to be set equal to the lower bound of the domain. The last four types enable you to force an activity to begin when its predecessor is finished. It is relatively easy to generate infeasible scenarios with the stricter versions, so you should use the stricter versions only if the weaker versions are not adequate for your problem.

Resource Constraints

The _RESOURCE_ and _QTY_ variables enable you to define resource constraints for activities. The _RESOURCE_ variable is a character variable that identifies the resource or resource pool. The _QTY_ variable is a numeric variable that identifies the number of units required. If the requirement is for a resource pool, you need to use the Resource data set to identify the pool members. See the section "Resource Data Set" on page 47 for more information.

For example, the following observations specify that activity A1 needs one unit of resource R1 and two units of resource R2:

```
_ACTIVITY_      _RESOURCE_      _QTY_

    A1              R1            1
    A1              R2            2
```

Variables in the ACTDATA= Data Set

Table 3.5 lists all the variables that are associated with the ACTDATA= data set and their interpretations by the CLP procedure. For each variable, the table also lists its type (C for character, N for numeric), its possible values, and its default value.

Table 3.5 Activity Data Set Variables

Name	Type	Description	Allowed Values	Default
ACTIVITY	C	Activity name		
DURATION	N	Duration	Nonnegative integers	0
SUCCESSOR	C	Successor name		
LAG	C	Lag type	FS, SS, FF, SF, FSE, SSE, FFE, SFE	FS
LAGDUR	N	Lag duration		0
ALIGNDATE	N	Alignment date		
ALIGNTYPE	C	Alignment type	SGE, SLE, SEQ, FGE, FLE, FEQ	
RESOURCE	C	Resource name		
QTY	N	Resource quantity	Nonnegative integers	1

Resource Data Set

The Resource data set is used in conjunction with the ACTDATA= data set to define resources, resource capacities, and alternate resources. The Resource data set contains at most four variables: _RESOURCE_, _CAPACITY_, _POOL_, and _SUBQTY_. The Resource data set is specified by using the RESDATA= option in the PROC CLP statement.

The _RESOURCE_ variable is a required character variable that defines resources. The _CAPACITY_ variable is a numeric variable that defines the capacity of the resource; it takes only nonnegative integer values. In the absence of alternate resources, the _RESOURCE_ and _CAPACITY_ variables are the only variables that you need in a data set to define resources and their capacities.

The following Resource data set defines resource R1 with capacity 2 and resource R2 with capacity 4:

```
_RESOURCE_      _CAPACITY_

    R1              2
    R2              4
```

Now suppose that you have an activity whose resource requirements can be satisfied by any one of a given set of resources. The Activity data set does not directly allow for a disjunctive specification. In order to do this you need to specify an abstract resource, referred to as a *resource pool*, in the _RESOURCE_ variable and use the Resource data set to identify the resources that can be substituted for this resource pool. You can do this using the _POOL_ and _SUBQTY_ variables. The _POOL_ variable is a character variable that identifies a resource pool to which the _RESOURCE_ variable belongs. The _SUBQTY_ variable is a numeric variable that identifies the number of units of _RESOURCE_ that can substitute for one unit of the resource pool. The _SUBQTY_ variable takes only nonnegative integer values. Each resource pool corresponds to as many observations in the Resource data set as there are members in the pool. A _RESOURCE_ can have membership in more than one resource pool. The resource and resource pool are distinct entities in the Resource data set;

that is, a _RESOURCE_ cannot have the same name as a _POOL_ in the Resource data set and vice versa.

For example, consider the following Activity data set:

```
Obs     _ACTIVITY_     _DURATION_      _RESOURCE_

 1           A             1              R1
 2           B             2              RP1
 3           C             1              RP2
```

and Resource data set:

```
Obs     _RESOURCE_     _CAPACITY_      _POOL_       _SUBQTY_

 1          R1             2            RP1            1
 2          R2             1            RP1            1
 3          R1             2            RP2            2
 4          R2             1            RP2            1
```

Activity A requires the resource R1. Activity B requires the resource RP1, which is identified as a resource pool in the Resource data set with members R1 and R2. Since the value of _SUBQTY_ is 1 for both resources, activity B can be satisfied with one unit of R1 or one unit of R2. Observations 3 and 4 in the Resource data set define resource pool RP2. Activity C requires resource pool RP2, which translates to requiring two units of R1 or one unit of R2 (since the value of _SUBQTY_ is 2 in observation 3 of the Resource data set). Resource substitution is not a sharable substitution; it is all or nothing. For example, if activity A requires two units of RP1 instead, the substitution is two units of R1 or two units of R2. The requirement cannot be satisfied using one unit of R1 and one unit of R2.

Variables in the RESDATA= Data Set

Table 3.6 lists all the variables that are associated with the RESDATA= data set and their interpretations by the CLP procedure. For each variable, the table also lists its type (C for character, N for numeric), its possible values, and its default value.

Table 3.6 *Resource Data Set Variables*

Name	Type	Description	Allowed Values	Default
RESOURCE	C	Resource name		
CAPACITY	N	Resource capacity	Nonnegative integers	1
POOL	C	Resource pool name		
SUBQTY	N	Number of units of resource that can substitute for one unit of the resource pool	Nonnegative integers	1

Schedule Data Set

In order to solve a scheduling type CSP, you need to specify one or more schedule-related output data sets by using one or more of the SCHEDULE=, SCHEDTIME=, or SCHEDRES= options in the PROC CLP statement.

The Schedule data set is specified with the SCHEDULE= option in the PROC CLP statement and is the only data set that contains both time and resource assignment information for each activity.

The SCHEDULE= data set always contains the following five variables: SOLUTION, ACTIVITY, DURATION, START, and FINISH. The SOLUTION variable gives the solution number to which each observation corresponds. The ACTIVITY variable identifies each activity. The DURATION variable gives the duration of the activity. The START and FINISH variables give the scheduled start and finish times for the activity. There is one observation that contains the time assignment information for each activity that corresponds to these variables.

If any resources have been specified, the data set contains three more variables: OBSTYPE, RE-SOURCE, and QTY. The value of the OBSTYPE variable indicates whether an observation represents time assignment information or resource assignment information. Observations that correspond to OBSTYPE="TIME" provide time assignment information, and observations that correspond to OBSTYPE="RESOURCE" provide resource assignment information. The RESOURCE variable and the QTY variable constitute the resource assignment information and identify the resource and quantity, respectively, of the resource that is assigned to each activity.

The values of RESOURCE and QTY are missing for time assignment observations, and the values of DURATION, START, and FINISH are missing for resource assignment observations.

If an Activity data set has been specified, the formats and labels for the _ACTIVITY_ and _DURA-TION_ variables carry over to the ACTIVITY and DURATION variables, respectively, in the Schedule data set.

In addition to or in lieu of the SCHEDULE= data set, there are two other schedule-related data sets that together represent a logical partitioning of the Schedule data set with no loss of data. The SCHEDTIME= data set contains the time assignment information, and the SCHEDRES= data set contains the resource assignment information.

Variables in the SCHEDULE= Data Set

Table 3.7 lists all the variables that are associated with the SCHEDULE= data set and their interpretations by the CLP procedure. For each variable, the table also lists its type (C for character, N for numeric), and its possible values.

Table 3.7 Schedule Data Set Variables

Name	Type	Description	Values
SOLUTION	N	Solution number	Positive integers
OBSTYPE	C	Observation type	TIME, RESOURCE
ACTIVITY	C	Activity name	
DURATION	N	Duration	Nonnegative integers, missing when OBSTYPE="RESOURCE"
START	N	Start time	Missing when OBSTYPE="RESOURCE"
FINISH	N	Finish time	Missing when OBSTYPE="RESOURCE"
RESOURCE	C	Resource name	Missing when OBSTYPE="TIME"
QTY	N	Resource quantity	Nonnegative integers, missing when OBSTYPE="TIME"

SCHEDRES= Data Set

The SCHEDRES= data set contains the resource assignments for each activity. There are four variables: SOLUTION, ACTIVITY, RESOURCE, and QTY, which are identical to the same variables in the SCHEDULE= data set. The observations correspond to the subset of observations in the SCHEDULE= data set with OBSTYPE="RESOURCE."

SCHEDTIME= Data Set

The SCHEDTIME= data set contains the time assignments for each activity. There are five variables: SOLUTION, ACTIVITY, DURATION, START, and FINISH, which are identical to the same variables in the SCHEDULE= data set. The observations correspond to the subset of observations in the SCHEDULE= data set with OBSTYPE="TIME."

Edge Finding

Edge-finding (EF) techniques are effective propagation techniques for resource capacity constraints that reason about the processing order of a set of activities that require a given resource or set of resources. Some of the typical ordering relationships that EF techniques can determine are whether an activity can, cannot, or must execute before (or after) a set of activities that require the same resource or set of resources. This in turn determines new time bounds on the start and finish times. Carlier and Pinson (1989) are responsible for some of the earliest work in this area, which resulted in solving MT10, a 10×10 job shop problem that had remained unsolved for over 20 years (Muth and Thompson 1963). Since then, there have been several variations and extensions of this work (Carlier and Pinson 1990; Applegate and Cook 1991; Nuijten 1994; Baptiste and Le Pape 1996).

The edge-finding consistency routines are invoked by specifying the EDGEFINDER= or EDGE= option in the SCHEDULE statement. Specifying EDGEFINDER=FIRST computes an upper bound on the activity finish time by detecting whether a given activity must be processed first from a set of activities that require the same resource or set of resources. Specifying EDGEFINDER=LAST computes a lower bound on the activity start time by detecting whether a given activity must be processed last from a set of activities that require the same resource or set of resources. Specifying EDGEFINDER=BOTH is equivalent to specifying both EDGEFINDER=FIRST and EDGEFINDER=LAST.

An extension of the edge-finding consistency routines is determining whether an activity *cannot* be the first to be processed or whether an activity *cannot* be the last to be processed from a given set of activities that require the same resource or set of resources. The NOTFIRST= or NF= option in the SCHEDULE statement determines whether an activity must not be the first to be processed. In similar fashion, the NOTLAST= or NL= option in the SCHEDULE statement determines whether an activity must not be the last to be processed.

Macro Variable _ORCLP_

The CLP procedure defines a macro variable named _ORCLP_. This variable contains a character string that indicates the status of the CLP procedure upon termination. The various terms of the macro variable are interpreted as follows.

STATUS

indicates the procedure status at termination. It can take one of the following values:

OK	The procedure terminated successfully.
DATA_ERROR	An input data error occurred.
IO_ERROR	A problem in reading or writing data occurred.
MEMORY_ERROR	Insufficient memory is allocated to the procedure.
SEMANTIC_ERROR	The use of semantic action is incorrect.
SYNTAX_ERROR	The use of syntax in incorrect.
ERROR	The status cannot be classified into any of the preceding categories.

If the procedure terminates normally or if an I/O error is detected while closing a data set, the following terms are added to the macro variable.

SOLUTION_STATUS

indicates the solution status at termination. It can take one of the following values:

ALL_SOLUTIONS	All solutions are found.
INFEASIBLE	The problem is infeasible.
SOLN_LIMIT_REACHED	The required number of solutions specified with the MAXSOLN= option is reached.
TIME_LIMIT_REACHED	The execution time limit specified with the MAXTIME= option is reached.
RESTART_LIMIT_REACHED	The number of restarts specified with RESTARTS= option is reached.
ABORT	The procedure is stopped by user before any other stop criteria is reached.

SOLUTIONS_FOUND

indicates the number of solutions that are found. This term is not applicable if SOLUTION_STATUS has a value of INFEASIBLE.

MIN_MAKESPAN

indicates the minimal makespan of the solutions that are found. The makespan is the maximum of the activity finish times or the completion time of the last job to leave the system. This term is applicable only to scheduling problems with at least one solution.

SOLUTION_TIME

indicates the time (in seconds) taken to solve the problem.

Examples: CLP Procedure

This section contains several examples that illustrate the capabilities of the different logical constraints and showcase a variety of problems that can be solved by the CLP procedure. The first seven examples feature a standard constraint satisfaction problem (CSP):

- "Example 3.1: Logic-Based Puzzles" illustrates the capabilities of the ALLDIFFERENT constraint in solving the popular logical puzzle, Sudoku. This example also contains a variant of Sudoku which illustrates the capabilities of the GCC constraint.

- "Example 3.2: Alphabet Blocks Problem" illustrates the use of the GCC constraint in solving the alphabet blocks problem, a popular combinatorial problem.

- "Example 3.3: Work-Shift Scheduling Problem" illustrates the capabilities of the ELEMENT constraint in modeling the cost information in a work-shift scheduling problem in order to find a minimum cost schedule.

- "Example 3.4: A Nonlinear Optimization Problem" illustrates how you can use the ELEMENT constraint to represent nonlinear functions and nonstandard variable domains, including noncontiguous domains.

- "Example 3.5: Car Painting Problem" involves limited sequencing of the cars in an assembly process in order to minimize the number of paint purgings and features the REIFY constraint.

- "Example 3.6: Scene Allocation Problem" schedules the shooting of different scenes of a movie in order to minimize production costs. This problem uses the GCC and LINEAR constraints.

- "Example 3.7: Car Sequencing Problem" relates to sequencing the cars on an assembly line with workstations for installing specific options subject to the demand constraints for each set of options and the capacity constraints of each workstation.

The next four examples feature scheduling CSPs and use the scheduling constraints in the CLP procedure:

- "Example 3.8: Round-Robin Problem" illustrates solving a single round robin tournament.

- "Example 3.9: Resource-Constrained Scheduling with Nonstandard Temporal Constraints" illustrates nonstandard precedence constraints in scheduling the construction of a bridge.

- "Example 3.10: Scheduling with Alternate Resources" illustrates a job scheduling problem with alternate resources. An optimal solution is determined by activating the edge-finding consistency techniques for this example.

- "Example 3.11: 10×10 Job Shop Scheduling Problem" illustrates a well-known 10×10 job shop scheduling problem and features edge-finding along with the edge-finding extensions 'not first' and 'not last' in order to determine optimality.

It is often possible to formulate a problem both as a standard CSP and also as a scheduling CSP. Depending on the nature of the constraints, it might even be more advantageous to formulate a scheduling problem as a standard CSP and vice versa:

- "Example 3.12: Scheduling a Major Basketball Conference" illustrates this concept by modeling the problem of scheduling a major basketball conference as a standard CSP. The ELEMENT constraint plays a key role in this particular example.

Example 3.1: Logic-Based Puzzles

There are many logic-based puzzles that can be formulated as CSPs. Two such instances are shown in this example.

Sudoku

Sudoku is a logic-based, combinatorial number-placement puzzle played on a partially filled 9×9 grid. The objective is to fill the grid with the digits 1 to 9, so that each column, each row, and each of the nine 3×3 blocks contain only one of each digit. Figure 3.5 shows an example of a Sudoku grid.

Figure 3.5 An Example of an Unsolved Sudoku Grid

This example illustrates the use of the ALLDIFFERENT constraint to solve the above Sudoku problem.

The data set indata contains the partially filled values for the grid and is used to create the set of macro variables C_{ij} ($i = 1 \ldots 9$, $j = 1 \ldots 9$), where C_{ij} is the value of cell (i, j) in the grid when specified, and missing otherwise.

```
data indata;
   input C1-C9;
   datalines;
. . 5 . . 7 . . 1
. 7 . . 9 . . 3 .
. . . 6 . . . . .
. . 3 . . 1 . . 5
. 9 . . 8 . . 2 .
1 . . 2 . . 4 . .
. . 2 . . 6 . . 9
. . . . 4 . . 8 .
8 . . 1 . . 5 . .
;
run;

%macro store_initial_values;
   /* store initial values into macro variable C_i_j */
```

```
   data _null_;
      set indata;
      array C{9};
      do j = 1 to 9;
         i = _N_;
         call symput(compress('C_'||put(i,best.)||'_'||put(j,best.)),
            put(C[j],best.));
      end;
   run;
%mend store_initial_values;


%store_initial_values;
```

Let the variable X_{ij} ($i = 1 \ldots 9, j = 1 \ldots 9$) represent the value of cell (i, j) in the grid. The domain of each of these variables is $[1, 9]$. Three sets of all-different constraints are used to set the required rules for each row, each column, and each of the 3×3 blocks. The constraint ALLDIFF($X_{i1} - X_{i9}$) forces all values in row i to be different, the constraint ALLDIFF($X_{1j} - X_{9j}$) forces all values in column j to be different, and the constraint ALLDIFF(X_{ij}) (($i = 1, 2, 3; j = 1, 2, 3$), ($i = 1, 2, 3; j = 4, 5, 6$), ..., ($i = 7, 8, 9; j = 7, 8, 9$)) forces all values in each block to be different.

The following statements solve the Sudoku puzzle:

```
%macro solve;
   proc clp out=outdata;

      /* Declare variables    */
      /* Nine row constraints */
      %do i = 1 %to 9;
         var (X_&i._1-X_&i._9) = [1,9];
         alldiff(X_&i._1-X_&i._9);
      %end;

      /* Nine column constraints */
      %do j = 1 %to 9;
         alldiff(
         %do i = 1 %to 9;
            X_&i._&j
         %end;
         );
      %end;

      /* Nine 3x3 block constraints */
      %do s = 0 %to 2;
         %do t = 0 %to 2;
            alldiff(
            %do i = 3*&s + 1 %to 3*&s + 3;
               %do j = 3*&t + 1 %to 3*&t + 3;
                  X_&i._&j
               %end;
            %end;
            );
         %end;
      %end;
```

```
        /* Initialize variables to cell values   */
        /* X_i_j = C_i_j if C_i_j is non-missing */
        %do i = 1 %to 9;
           %do j = 1 %to 9;
              %if &&C_&i._&j ne . %then %do;
                 lincon X_&i._&j = &&C_&i._&j;
              %end;
           %end;
        %end;

     run;
     %put &_ORCLP_;
  %mend solve;

  %solve
```

Output 3.1.1 shows the solution.

Output 3.1.1 Solution of the Sudoku Grid

9	8	5	3	2	7	6	4	1
6	7	1	5	9	4	2	3	8
3	2	4	6	1	8	9	5	7
2	4	3	7	6	1	8	9	5
5	9	7	4	8	3	1	2	6
1	6	8	2	5	9	4	7	3
4	5	2	8	3	6	7	1	9
7	1	6	9	4	5	3	8	2
8	3	9	1	7	2	5	6	4

The basic structure of the classical Sudoku problem can easily be extended to formulate more complex puzzles. One such example is the Pi Day Sudoku puzzle.

Pi Day Sudoku

Pi Day is a celebration of the number π that occurs every March 14. In honor of Pi Day, Brainfreeze Puzzles (Riley and Taalman 2008) celebrates this day with a special 12×12 grid Sudoku puzzle. The 2008 Pi Day Sudoku puzzle is shown in Figure 3.6.

Figure 3.6 Pi Day Sudoku 2008

The rules of this puzzle are a little different from the standard Sudoku. First, the blocks in this puzzle are jigsaw regions rather than 3×3 blocks. Each jigsaw region consists of 12 contiguous cells. Second, the first 12 digits of π are used instead of the digits 1–9. Each row, column, and jigsaw region contains the first 12 digits of π (314159265358) in some order. In particular, there are two 1s, two 3s, three 5s, no 7s, and one 2, 4, 6, 8, and 9.

The data set raw contains the partially filled values for the grid and, similar to the Sudoku problem, is used to create the set of macro variables C_{ij} ($i = 1, \ldots, 12$, $j = 1, \ldots, 12$) where C_{ij} is the value of cell (i, j) in the grid when specified, and missing otherwise.

```
data raw;
   input C1-C12;
   datalines;
3  .  .  1  5  4  .  .  1  .  9  5
.  1  .  .  3  .  .  .  .  1  3  6
.  .  4  .  .  3  .  8  .  .  2  .
5  .  .  1  .  .  9  2  5  .  .  1
.  9  .  .  5  .  .  5  .  .  .  .
5  8  1  .  .  9  .  .  3  .  6  .
.  5  .  8  .  .  2  .  .  5  5  3
.  .  .  .  5  .  .  6  .  .  1  .
2  .  .  5  1  5  .  .  5  .  .  9
.  6  .  .  4  .  1  .  .  3  .  .
1  5  1  .  .  .  .  5  .  .  5  .
5  5  .  4  .  .  3  1  6  .  .  8
;
run;
```

```
%macro cdata;
   /* store each pre-filled value into macro variable C_i_j */
   data _null_;
      set raw;
      array C{12};
      do j = 1 to 12;
         i = _N_;
         call symput(compress('C_'||put(i,best.)||'_'||put(j,best.)),
            put(C[j],best.));
      end;
   run;
%mend cdata;
%cdata;
```

As in the Sudoku problem, let the variable X_{ij} represent the value of the cell that corresponds to row i and column j. The domain of each of these variables is $[1, 9]$.

For each row, column, and jigsaw region, a GCC statement is specified to enforce the condition that it contains exactly the first twelve digits of π.

In particular, the variables in row r, $r = 1, \ldots, 12$ are $X_{r1}, \ldots, X_{r12}$. The SAS macro %CONS_ROW(R) enforces the GCC constraint that row r contains exactly two 1s, two 3s, three 5s, no 7s, and one of each of the other values:

```
%macro cons_row(r);
   /* Row r must contain two 1's, two 3's, three 5's, no 7's, */
   /* and one for each of other values from 1 to 9.           */
   gcc(X_&r._1-X_&r._12) =
      ( (1, 2, 2) (3, 2, 2) (5, 3, 3) (7, 0, 0) DL=1 DU=1 );
%mend cons_row;
```

The variables in column c are $X_{1c}, \ldots, X_{12c}$. The SAS macro %CONS_COL(C) enforces a similar GCC constraint for each column c.

```
%macro cons_col(c);
   /* Column c must contain two 1's, two 3's, three 5's,    */
   /* no 7's, and one for each of other values from 1 to 9. */
   gcc( %do r = 1 %to 12;
      X_&r._&c.
      %end;
      ) = ((1, 2, 2) (3, 2, 2) (5, 3, 3) (7, 0, 0) DL=1 DU=1);
%mend cons_col;
```

Generalizing this concept further, the SAS macro %CONS_REGION(VARS) enforces the GCC constraint for the jigsaw region that is defined by the macro variable VARS.

```
%macro cons_region(vars);
   /* Jigsaw region that contains &vars must contain two 1's, */
   /* two 3's, three 5's, no 7's, and one for each of other   */
   /* values from 1 to 9.                                     */
   gcc(&vars.) = ((1, 2, 2) (3, 2, 2) (5, 3, 3) (7, 0, 0) DL=1 DU=1);
%mend cons_region;
```

The following SAS statements incorporate the preceding macros to define the GCC constraints in order to find all solutions of the Pi Day Sudoku 2008 puzzle:

```
%macro pds(solns=allsolns,varsel=MINR,maxt=900);

   proc clp out=pdsout &solns
           varselect=&varsel /* Variable selection strategy */
           maxtime=&maxt;     /* Time limit                 */

       /* Variable X_i_j represents the grid of ith row and jth column. */
       var (
          %do i = 1 %to 12;
             X_&i._1 - X_&i._12
          %end;
          ) = [1,9];

       /* X_i_j = C_i_j if C_i_j is non-missing */
       %do i = 1 %to 12;
          %do j = 1 %to 12;
             %if &&C_&i._&j ne . %then %do;
                lincon X_&i._&j = &&C_&i._&j;
             %end;
          %end;
       %end;

       /* 12 Row constraints: */
       %do r = 1 %to 12;
          %cons_row(&r);
       %end;

       /* 12 Column constraints: */
       %do c = 1 %to 12;
          %cons_col(&c);
       %end;

       /* 12 Jigsaw region constraints: */
       /* Each jigsaw region is defined by the macro variable &vars. */

       /* Region 1: */
       %let vars = X_1_1 - X_1_3 X_2_1 - X_2_3
                   X_3_1 X_3_2 X_4_1 X_4_2 X_5_1 X_5_2;
       %cons_region(&vars.);

       /* Region 2: */
       %let vars = X_1_4 - X_1_9 X_2_4 - X_2_9;
       %cons_region(&vars.);

       /* Region 3: */
       %let vars = X_1_10 - X_1_12 X_2_10 - X_2_12
                   X_3_11 X_3_12 X_4_11 X_4_12 X_5_11 X_5_12;
       %cons_region(&vars.);

       /* Region 4: */
       %let vars = X_3_3 - X_3_6 X_4_3 - X_4_6 X_5_3 - X_5_6;
       %cons_region(&vars.);
```

```
        /* Region 5: */
        %let vars = X_3_7 - X_3_10 X_4_7 - X_4_10 X_5_7 - X_5_10;
        %cons_region(&vars.);

        /* Region 6: */
        %let vars = X_6_1 - X_6_3 X_7_1 - X_7_3
                    X_8_1 - X_8_3 X_9_1 - X_9_3;
        %cons_region(&vars.);

        /* Region 7: */
        %let vars = X_6_4 X_6_5 X_7_4 X_7_5 X_8_4 X_8_5
                    X_9_4 X_9_5 X_10_4 X_10_5 X_11_4 X_11_5;
        %cons_region(&vars.);

        /* Region 8: */
        %let vars = X_6_6 X_6_7 X_7_6 X_7_7 X_8_6 X_8_7
                    X_9_6 X_9_7 X_10_6 X_10_7 X_11_6 X_11_7;
        %cons_region(&vars.);

        /* Region 9: */
        %let vars = X_6_8 X_6_9 X_7_8 X_7_9 X_8_8 X_8_9
                    X_9_8 X_9_9 X_10_8 X_10_9 X_11_8 X_11_9;
        %cons_region(&vars.);

        /* Region 10: */
        %let vars = X_6_10 - X_6_12 X_7_10 - X_7_12
                    X_8_10 - X_8_12 X_9_10 - X_9_12;
        %cons_region(&vars.);

        /* Region 11: */
        %let vars = X_10_1 - X_10_3 X_11_1 - X_11_3 X_12_1 - X_12_6;
        %cons_region(&vars.);

        /* Region 12: */
        %let vars = X_10_10 - X_10_12 X_11_10 - X_11_12 X_12_7 - X_12_12;
        %cons_region(&vars.);
    run;
    %put &_ORCLP_;

%mend pds;

%pds;
```

The only solution of the 2008 Pi Day Sudoku puzzle is shown in Output 3.1.2.

Output 3.1.2 Solution to Pi Day Sudoku 2008

```
                              Pi Day Sudoku 2008

  Obs     C1     C2     C3     C4     C5     C6     C7     C8     C9    C10    C11    C12

   1       3      2      5      1      5      4      6      3      1      8      9      5
   2       4      1      5      2      3      8      5      9      5      1      3      6
   3       6      1      4      5      9      3      5      8      3      1      2      5
   4       5      3      3      1      8      5      9      2      5      6      4      1
   5       8      9      2      6      5      1      1      5      4      3      3      5
   6       5      8      1      5      2      9      4      3      3      5      6      1
   7       1      5      3      8      1      6      2      4      9      5      5      3
   8       9      4      5      3      5      1      5      6      8      2      1      3
   9       2      3      6      5      1      5      3      1      5      4      8      9
  10       3      6      8      9      4      5      1      5      1      3      5      2
  11       1      5      1      3      6      3      8      5      2      9      5      4
  12       5      5      9      4      3      2      3      1      6      5      1      8
```

The corresponding completed grid is shown in Figure 3.7.

Figure 3.7 Solution to Pi Day Sudoku 2008

Example 3.2: Alphabet Blocks Problem

This example illustrates usage of the global cardinality constraint (GCC). The alphabet blocks problem consists of finding an arrangement of letters on four alphabet blocks. Each alphabet block has a single letter on each of its six sides. Collectively, the four blocks contain every letter of the alphabet except Q and Z. By arranging the blocks in various ways, the following words should be spelled out: BAKE, ONYX, ECHO, OVAL, GIRD, SMUG, JUMP, TORN, LUCK, VINY, LUSH, and WRAP.

You can formulate this problem as a CSP by representing each of the 24 letters with an integer variable. The domain of each variable is the set $\{1, 2, 3, 4\}$ that represents block1 through block4. The assignment '$A = 1$' indicates that the letter 'A' is on a side of block1. Each block has six sides; hence each value v in $\{1, 2, 3, 4\}$ has to be assigned to exactly six variables so that each side of a block has a letter on it. This restriction can be formulated as a global cardinality constraint over all 24 variables with common lower and upper bounds set equal to six.

Moreover, in order to spell all of the words listed previously, the four letters in each of the 12 words have to be on different blocks. Another GCC statement that specifies 12 global cardinality constraints is used to enforce these conditions. You can also formulate these restrictions with 12 all-different constraints. Finally, four linear constraints (as specified with LINCON statements) are used to break the symmetries that blocks are interchangeable. These constraints preset the blocks that contain the letters 'B', 'A', 'K', and 'E' as block1, block2, block3, and block4, respectively.

The complete representation of the problem is as follows:

```
proc clp out=out;
   /* Each letter except Q and Z is represented with a variable. */
   /* The domain of each variable is the set of 4 blocks,         */
   /*    or {1, 2, 3, 4} for short.                               */
   var (A B C D E F G H I J K L M N O P R S T U V W X Y) = [1,4];

   /* There are exactly 6 letters on each alphabet block */
   gcc (A B C D E F G H I J K L M N O P R S T U V W X Y) = (
                                        (1, 6, 6)
                                        (2, 6, 6)
                                        (3, 6, 6)
                                        (4, 6, 6) );

   /* Note 1: Since lv=uv=6 for all v=1,...,4; the above global
              cardinality constraint can also specified as:
   gcc (A B C D E F G H I J K L M N O P R S T U V W X Y) =(DL=6 DU=6);
   */
   /* The letters in each word must be on different blocks. */
   gcc (B A K E) = (DL=0 DU=1)
       (O N Y X) = (DL=0 DU=1)
       (E C H O) = (DL=0 DU=1)
       (O V A L) = (DL=0 DU=1)
       (G I R D) = (DL=0 DU=1)
       (S M U G) = (DL=0 DU=1)
       (J U M P) = (DL=0 DU=1)
       (T O R N) = (DL=0 DU=1)
       (L U C K) = (DL=0 DU=1)
       (V I N Y) = (DL=0 DU=1)
       (L U S H) = (DL=0 DU=1)
       (W R A P) = (DL=0 DU=1);

   /* Note 2: These restrictions can also be enforced by ALLDIFF constraints:
         alldiff (B A K E) (O N Y X) (E C H O) (O V A L)
                 (G I R D) (S M U G) (J U M P) (T O R N)
                 (L U C K) (V I N Y) (L U S H) (W R A P);
   */
```

```
      /* Breaking the symmetry that blocks can be interchanged by setting
         the block that contains the letter B as block1, the block that
         contains the letter A as block2, etc. */
      lincon B = 1;
      lincon A = 2;
      lincon K = 3;
      lincon E = 4;

   run;
```

The solution to this problem is shown in Output 3.2.1.

Output 3.2.1 Solution to Alphabet Blocks Problem

```
                  Solution to Alphabet Blocks Problem

        Block    Side1    Side2    Side3    Side4    Side5    Side6

          1        B        F        I        O        U        W
          2        A        C        D        J        N        S
          3        H        K        M        R        V        X
          4        E        G        L        P        T        Y
```

Example 3.3: Work-Shift Scheduling Problem

This example illustrates the use of the GCC constraint in finding a feasible solution to a work-shift scheduling problem and then using the element constraint to incorporate cost information in order to find a minimum cost schedule.

Six workers (Alan, Bob, John, Mike, Scott, and Ted) are to be assigned to three working shifts. The first shift needs at least one and at most four people; the second shift needs at least two and at most three people; and the third shift needs exactly two people. Alan does not work on the first shift; Bob works only on the third shift. The others can work any shift. The objective is to find a feasible assignment for this problem.

You can model the minimum and maximum shift requirements with a GCC constraint and formulate the problem as a standard CSP. The variables W1–W6 identify the shift to be assigned to each of the six workers: Alan, Bob, John, Mike, Scott, and Ted.

```
proc clp out=clpout;
   /*  Six workers (Alan, Bob, John, Mike, Scott and Ted)
       are to be assigned to 3 working shifts.           */
   var (W1-W6) = [1,3];

   /* The first shift needs at least 1 and at most 4 people;
      the second shift needs at least 2 and at most 3 people;
      and the third shift needs exactly 2 people. */
   gcc (W1-W6) = ( ( 1, 1, 4) ( 2, 2, 3) ( 3, 2, 2) );
```

```
        /* Alan doesn't work on the first shift. */
        lincon W1 <> 1;

        /* Bob works only on the third shift. */
        lincon W2 = 3;
    run;
```

The resulting assignment is shown in Output 3.3.1.

Output 3.3.1 Solution to Work-Shift Scheduling Problem

```
                Solution to Work-Shift Scheduling Problem

            Obs     W1     W2     W3     W4     W5     W6

             1      2      3      1      1      2      3
```

A Gantt chart of the corresponding schedule is displayed in Output 3.3.2.

Output 3.3.2 Work-Shift Schedule
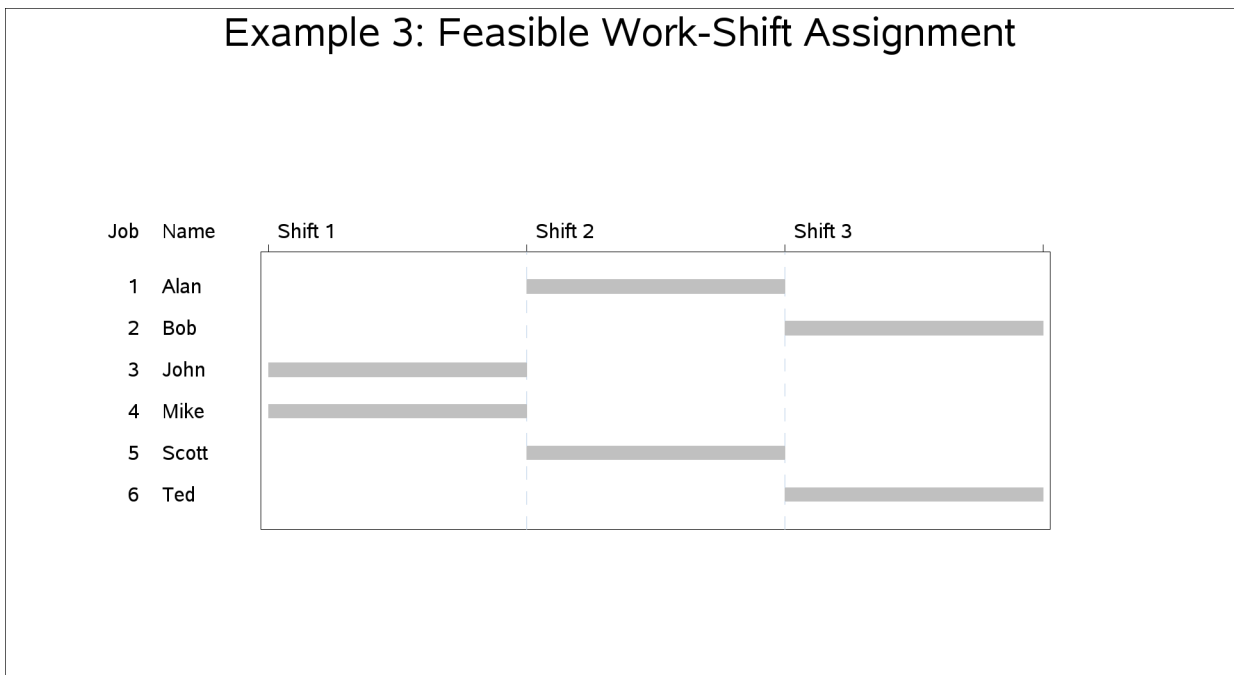
Now suppose that every work-shift assignment has a cost associated with it and that the objective of interest is to determine the schedule with minimum cost.

The costs of assigning the workers to the different shifts are given in Table 3.8. A dash "-" in position (i, j) indicates that worker i can not work on shift j.

Table 3.8 Costs of Assigning Workers to Shifts

	Shift 1	Shift 2	Shift 3
Alan	-	12	10
Bob	-	-	6
John	16	8	12
Mike	10	6	8
Scott	6	6	8
Ted	12	4	4

Based on the cost structure in Table 3.8, the schedule derived previously has a cost of 54. The objective now is to determine the optimal schedule—one that results in the minimum cost.

Let the variable C_i represent the cost of assigning worker i to a shift. This variable is shift-dependent and is given a high value (for example, 100) if the worker cannot be assigned to a shift. The costs can also be interpreted as preferences if desired. You can use an element constraint to associate the cost C_i with the shift assignment for each worker. For example, C_1, the cost of assigning Alan to a shift, can be determined by the constraint ELEMENT(W_1, (100, 12, 10), C_1).

By adding a linear constraint $\sum_{i=1}^{n} C_i \leq obj$, you can limit the solutions to feasible schedules that cost no more than *obj*.

You can then create a SAS macro %CALLCLP with *obj* as a parameter that can be called iteratively from a search routine to find an optimal solution. The SAS macro %MINCOST(*lb,ub*) uses a bisection search to find the minimum cost schedule among all schedules that cost between *lb* and *ub*. Although a value of *ub* = 100 is used in this example, it would suffice to use *ub* = 54, the cost of the feasible schedule determined earlier.

```
%macro callclp(obj);
   %put The objective value is: &obj..;
   proc clp out=clpout;
      /* Six workers (Alan, Bob, John, Mike, Scott and Ted)
         are to be assigned to 3 working shifts.   */
      var (W1-W6) = [1,3];
      var (C1-C6) = [1,100];

      /* The first shift needs at least 1 and at most 4 people;
         the second shift needs at least 2 and at most 3 people;
         and the third shift needs exactly 2 people. */
      gcc (W1-W6) = ( ( 1, 1, 4) ( 2, 2, 3) ( 3, 2, 2) );

      /* Alan doesn't work on the first shift. */
      lincon W1 <> 1;

      /* Bob works only on the third shift. */
      lincon W2 = 3;
```

```
        /* Specify the costs of assigning the workers to the shifts.
           Use 100 (a large number) to indicate an assignment
           that is not possible.*/
        element (W1, (100,  12, 10), C1);
        element (W2, (100, 100,  6), C2);
        element (W3, ( 16,   8, 12), C3);
        element (W4, ( 10,   6,  8), C4);
        element (W5, (  6,   6,  8), C5);
        element (W6, ( 12,   4,  4), C6);

        /* The total cost should be no more than the given objective value. */
        lincon C1 + C2 + C3 + C4 + C5 + C6 <= &obj;
    run;

    /* when a solution is found,                          */
    /* &_ORCLP_ contains the string SOLUTIONS_FOUND   */
    %if %index(&_ORCLP_,  SOLUTIONS_FOUND) %then %let clpreturn=SUCCESSFUL;
%mend;

/* Bisection search method to determine the optimal objective value  */
%macro mincost(lb, ub);
    %do %while (&lb<&ub-1);
        %put Currently lb=&lb, ub=&ub..;
        %let newobj=%eval((&lb+&ub)/2);
        %let clpreturn=NOTFOUND;
        %callclp(&newobj);
        %if &clpreturn=SUCCESSFUL %then %let ub=&newobj;
        %else %let lb=&newobj;
    %end;
    %callclp(&ub);
    %put Minimum possible objective value within given range is &ub.;
    %put Any value less than &lb makes the problem infeasible. ;
    proc print;
    run;
%mend;

/* Find the minimum objective value between 1 and 100. */
%mincost(lb=1, ub=100);
```

The cost of the optimal schedule, which corresponds to the solution shown in the following output, is 40.

```
           Solution to Optimal Work-Shift Scheduling Problem

  Obs    W1    W2    W3    W4    W5    W6    C1    C2    C3    C4    C5    C6

   1      3     3     2     2     1     2    10     6     8     6     6     4
```

The minimum cost schedule is displayed in the Gantt chart in Output 3.3.3.

Output 3.3.3 Work-Shift Schedule with Minimum Cost

Example 3: Optimal Work-Shift Assignment
Minimum Cost Schedule: $40
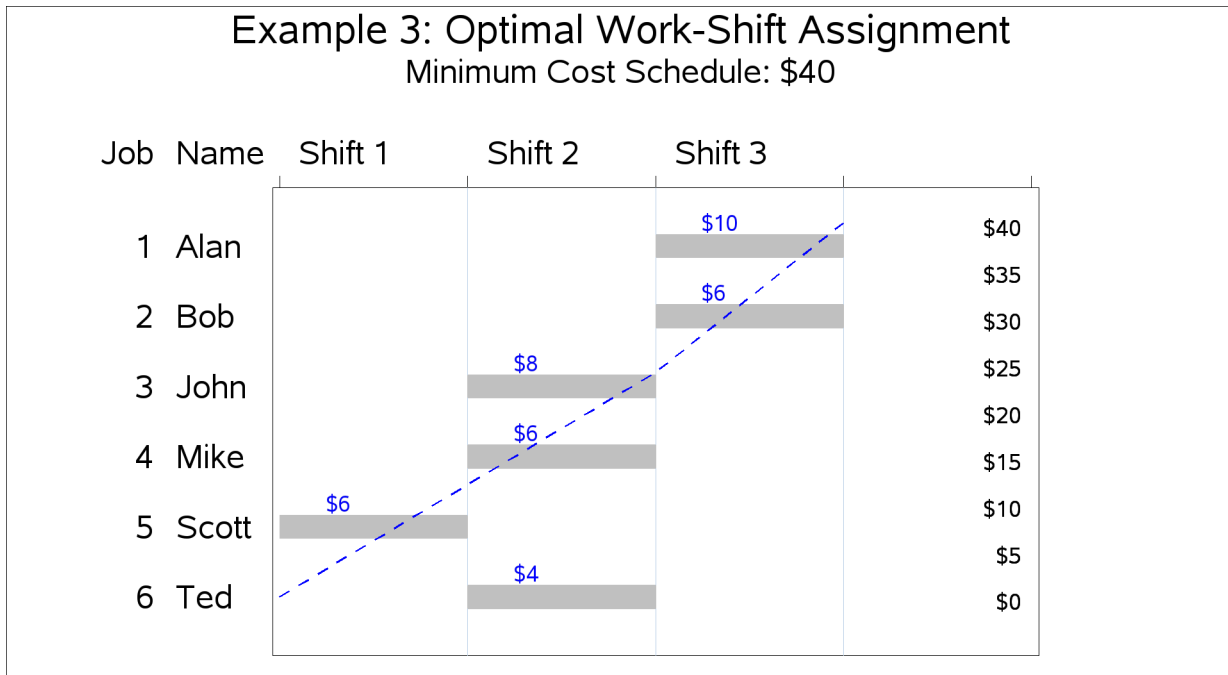
Example 3.4: A Nonlinear Optimization Problem

This example illustrates how you can use the element constraint to represent almost any function between two variables in addition to representing nonstandard domains. Consider the following nonlinear optimization problem:

$$\text{maximize } f(x) = x_1^3 + 5x_2 - 2^{x_3}$$

$$\text{subject to } \begin{cases} x_1 & - & .5x_2 & + & x_3^2 & \leq & 50 \\ \text{mod}(x_1, 4) & + & .25x_2 & & & \geq & 1.5 \end{cases}$$

$$x_1 : \text{integers in } [-5, 5], \ x_2 : \text{odd integers in } [-5, 9], \ x_3 : \text{integers in } [1, 10].$$

You can use the CLP procedure to solve this problem by introducing four artificial variables y_1–y_4 to represent each of the nonlinear terms. Let $y_1 = x_1^3$, $y_2 = 2^{x_3}$, $y_3 = x_3^2$, and $y_4 = \text{mod}(x_1, 4)$. Since the domains of x_1 and x_2 are not consecutive integers that start from 1, you can use element constraints to represent their domains by using index variables z_1 and z_2, respectively. For example, either of the following two ELEMENT constraints specifies that the domain of x_2 is the set of odd integers in $[-5, 9]$:

```
element(z2,(-5,-3,-1,1,3,5,7,9),x2)
element(z2,(-5 to 9 by 2),x2)
```

Any functional dependencies on x_1 or x_2 can now be defined using z_1 or z_2, respectively, as the index variable in an element constraint. Since the domain of x_3 is $[1, 10]$, you can directly use x_3 as the index variable in an element constraint to define dependencies on x_3.

For example, the following constraint specifies the function $y_1 = x_1^3$, $x_1 \in [-5, 5]$

```
element(z1,(-125,-64,-27,-8,-1,0,1,8,27,64,125),y1)
```

The nonlinear optimization problem can now be reduced to a set of element and linear constraints. By expressing the objective function as a linear constraint $f(x) \geq obj$ where *obj* is a supplied parameter, you can find a solution with objective value of at least *obj*.

In a manner similar to that of Example 3.3, you can create a SAS macro %CALLCLP with parameter *obj* that can be called iteratively from a search method such as %FINDMAX to determine the optimal value of the objective function, as follows:

```
%macro callclp(obj);
%put The objective value is: &obj..;
proc clp out=clpout;
   var x1=[-5, 5] x2=[-5, 9] x3=[1, 10] (y1-y4) (z1-z2);

   /* Use element constraint to represent non-contiguous domains */
   /* and nonlinear functions.                                   */
   element

      /* Domain of x1 is [-5,5] */
      (z1, ( -5 to 5), x1)

      /* Functional Dependencies on x1 */
      /* y1 = x1^3 */
      (z1, (-125, -64, -27, -8, -1, 0, 1, 8, 27, 64, 125), y1)
      /* y4 = mod(x1, 4) */
      (z1, ( -1,  0,  -3,  -2,  -1, 0, 1, 2, 3, 0, 1), y4)

      /* Domain of x2 is the set of odd numbers in [-5, 9] */
      (z2, (-5 to 9 by 2), x2)

      /* Functional Dependencies on x3 */
      /* y2 = 2^x3 */
      (x3, (2, 4, 8, 16, 32, 64, 128, 256, 512, 1024), y2)
      /* y3 = x3^2 */
      (x3, (1, 4, 9, 16, 25, 36, 49, 64, 81, 100), y3);

   lincon
      /* Objective function: x1^3 + 5 * x2 - 2^x3 */
      y1 + 5 * x2 - y2 >= &obj,

      /* x1 - .5 * x2 + x3^2 <=50 */
      x1 - .5 * x2 + y3 <= 50,

      /* mod(x1, 4) + .25 * x2 >= 1.5 */
      y4 + .25 * x2 >= 1.5;
run;
```

```
/* when a solution is found,                         */
/* &_ORCLP_ contains the string SOLUTIONS_FOUND      */
%if %index(&_ORCLP_,  SOLUTIONS_FOUND) %then %let clpreturn=SUCCESSFUL;
%mend;


/* Bisection search to determine the optimal objective value. */
%macro findmax(lb, ub);
   %do %while (&lb<&ub-1);
      %put Currently lb=&lb, ub=&ub..;
      %let newobj=%eval((&lb+&ub)/2);
      %let clpreturn=NOTFOUND;
      %callclp(&newobj);
      %if &clpreturn=SUCCESSFUL %then %let lb=&newobj;
      %else %let ub=&newobj;
   %end;
   %callclp(&lb);
   %put Maximum possible objective value within given range is &lb.;
   %put Any value greater than &lb makes the problem infeasible.;
   proc print;
   run;
%mend;


/*Find the maximum objective value between -200 and 200. */
%findmax(lb=-200, ub=200);
```

Output 3.4.1 shows the solution that corresponds to the optimal objective value of 168.

Output 3.4.1 Nonlinear Optimization Problem Solution

Obs	x1	x2	x3	y1	y2	y3	y4	z1	z2
1	5	9	1	125	2	1	1	11	8

Example 3.5: Car Painting Problem

The car painting process is an important part of the automobile manufacturing industry. Purging (the act of changing colors in the assembly process) is expensive due to the added cost of wasted paint and solvents involved with each color change in addtion to the extra time required for the purging process. The objective in the car painting problem is to sequence the cars in the assembly in order to minimize paint changeover (Sokol 2002; Trick 2004).

There are 10 cars in a sequence. The order for assembly is 1, 2, ..., 10. A car must be painted within three positions of its assembly order. For instance, car 5 can be painted in positions 2 through 8 inclusive. Cars 1, 5, and 9 are red; 2, 6, and 10 are blue; 3 and 7 green; and 4 and 8 are yellow. The initial sequence 1, 2, ..., 10 corresponds to the color pattern RBGYRBGYRB and has 9 purgings. The objective is to find a solution that minimizes the number of purgings.

This problem can be formulated as a CSP as follows. The variables S_i and C_i represent the ID and color, respectively, of the car in slot i. An element constraint relates the car ID to its color. An all-different constraint ensures that every slot is associated with a unique car ID. Two linear constraints represent the constraint that a car must be painted within three positions of its assembly order. The binary variable P_i indicates whether a paint purge takes place after the car in slot i is painted. Finally, a linear constraint is used to limit the total number of purgings to the required number.

The following %CAR_PAINTING macro determines all feasible solutions for a given number of purgings, which is specified as a parameter to the macro:

```
%macro car_painting(purgings);

   proc clp out=car_ds findall;

      %do i = 1 %to 10;
         var S&i = [1, 10]; /* which car is in slot &i.*/
         var C&i = [1, 4]; /* which color the car in slot &i is.*/
         /* Red=1; Blue=2; Green=3; Yellow=4 */
         element (S&i, (1, 2, 3, 4, 1, 2, 3, 4, 1, 2), C&i);
      %end;

      /* A car can be painted only once. */
      alldiff (S1-S10);

      /* A car must be painted within 3 positions of its assembly order. */
      %do i = 1 %to 10;
         lincon S&i-&i>=-3;
         lincon S&i-&i<=3;
      %end;

      %do i = 1 %to 9;
         var P&i = [0, 1]; /* Whether there is a purge after slot &i*/
         reify P&i: (C&i <> C%eval(&i+1));
      %end;

      /* Calculate the number of purgings. */
      lincon 0
      %do i = 1 %to 9;
         + P&i
      %end;
      <=&purgings ;
   run;

%mend;
%car_painting(5)
```

The problem is infeasible for four purgings. The CLP procedure finds 87 possible solutions for the five purgings problem. The solutions are sorted by the total distance all cars are moved in the sequencing, which ranges from 12 to 22 slots. The first 15 solutions are displayed in the Gantt chart in Output 3.5.1. Each row represents a solution, and each color transition represents a paint purge.

Output 3.5.1 Car Painting Schedule with Five Purgings

Car Painting Problem
Solutions with Five Purgings

Solution	Distance	1	2	3	4	5	6	7	8	9	10
1	12	Car 1: red	Car 5: red	Car 3: green	Car 7: green	Car 2: blue	Car 6: blue	Car 4: yellow	Car 8: yellow	Car 9: red	Car 10: blue
2	12	Car 1: red	Car 5: red	Car 2: blue	Car 6: blue	Car 3: green	Car 7: green	Car 4: yellow	Car 8: yellow	Car 9: red	Car 10: blue
3	12	Car 1: red	Car 2: blue	Car 6: blue	Car 3: green	Car 7: green	Car 4: yellow	Car 8: yellow	Car 5: red	Car 9: red	Car 10: blue
4	12	Car 1: red	Car 2: blue	Car 6: blue	Car 4: yellow	Car 8: yellow	Car 3: green	Car 7: green	Car 5: red	Car 9: red	Car 10: blue
5	12	Car 1: red	Car 2: blue	Car 3: green	Car 7: green	Car 5: red	Car 9: red	Car 4: yellow	Car 8: yellow	Car 6: blue	Car 10: blue
6	12	Car 1: red	Car 2: blue	Car 3: green	Car 7: green	Car 4: yellow	Car 8: yellow	Car 5: red	Car 9: red	Car 6: blue	Car 10: blue
7	14	Car 1: red	Car 5: red	Car 3: green	Car 7: green	Car 2: blue	Car 6: blue	Car 4: yellow	Car 8: yellow	Car 10: blue	Car 9: red
8	14	Car 1: red	Car 5: red	Car 2: blue	Car 6: blue	Car 7: green	Car 3: green	Car 4: yellow	Car 8: yellow	Car 9: red	Car 10: blue
9	14	Car 1: red	Car 5: red	Car 6: blue	Car 2: blue	Car 3: green	Car 7: green	Car 4: yellow	Car 8: yellow	Car 9: red	Car 10: blue
10	14	Car 1: red	Car 5: red	Car 2: blue	Car 6: blue	Car 3: green	Car 7: green	Car 4: yellow	Car 8: yellow	Car 10: blue	Car 9: red
11	14	Car 1: red	Car 5: red	Car 2: blue	Car 3: green	Car 7: green	Car 4: yellow	Car 8: yellow	Car 6: blue	Car 10: blue	Car 9: red
12	14	Car 1: red	Car 5: red	Car 2: blue	Car 4: yellow	Car 8: yellow	Car 3: green	Car 7: green	Car 6: blue	Car 10: blue	Car 9: red
13	14	Car 1: red	Car 5: red	Car 2: blue	Car 3: green	Car 7: green	Car 4: yellow	Car 8: yellow	Car 9: red	Car 6: blue	Car 10: blue
14	14	Car 1: red	Car 5: red	Car 2: blue	Car 4: yellow	Car 8: yellow	Car 3: green	Car 7: green	Car 9: red	Car 6: blue	Car 10: blue
15	14	Car 1: red	Car 2: blue	Car 6: blue	Car 7: green	Car 3: green	Car 4: yellow	Car 8: yellow	Car 5: red	Car 9: red	Car 10: blue

Example 3.6: Scene Allocation Problem

The scene allocation problem consists of deciding when to shoot each scene of a movie in order to minimize the total production cost (Van Hentenryck 2002). Each scene involves a number of actors, and at most five scenes a day can be shot. All actors who appear in a scene must be present on the day the scene is shot. Each actor has a daily rate for each day spent in the studio, regardless of the number of scenes in which he or she appears on that day. The goal is to minimize the production costs of the studio.

The actor names, daily fees, and the scenes they appear are given in the SCENE data set shown in Output 3.6.1. The variables S_Var1, ..., S_Var9 indicate the scenes in which the actor appears. For example, the first observation indicates that Patt's daily fee is 26481 and that Patt appears in scenes 2, 5, 7, 10, 11, 13, 15, and 17.

Output 3.6.1 The Scene Data Set

Obs	Num ber	Actor	Daily Fee	S_Var1	S_Var2	S_Var3	S_Var4	S_Var5	S_Var6	S_Var7	S_Var8	S_Var9
1	1	Patt	26481	2	5	7	10	11	13	15	17	.
2	2	Casta	25043	4	7	9	10	13	16	19	.	.
3	3	Scolaro	30310	3	6	9	10	14	16	17	18	.
4	4	Murphy	4085	2	8	12	13	15	.	.	.	.
5	5	Brown	7562	2	3	12	17	.	.	.	.	.
6	6	Hacket	9381	1	2	12	13	18	.	.	.	.
7	7	Anderson	8770	5	6	14	.	.	.	.	.	.
8	8	McDougal	5788	3	5	6	9	10	12	15	16	18
9	9	Mercer	7423	3	4	5	8	9	16	.	.	.
10	10	Spring	3303	5	6	.	.	.	.	.	.	.
11	11	Thompson	9593	6	9	12	15	18	.	.	.	.

There are nineteen scenes and at most five scenes can be filmed in one day, so at least four days are needed to schedule all the scenes ($\lceil \frac{19}{5} \rceil = 4$). Let Sj_k be a binary variable that equals 1 if scene j is shot on day k. Let Ai_k be another binary variable that equals 1 if actor i is present on day k.

The following %SCENECLP macro invokes the CLP procedure. It includes two sets of GCC constraints: one to make sure each scene is shot exactly once, and one to limit the number of scenes shot per day to be at least four and at most five. There are also three sets of LINCON constraints: one to indicate that an actor must be present if any of his or her scenes are shot that day, one for breaking symmetries to improve efficiency, and one to limit the total production cost to be less than or equal to the given objective.

```
%macro sceneclp(obj);
   %put The objective value is: &obj..;

   proc clp out=out varselect=maxc;
      /* Declare variables. */
      %do k=1 %to 4; /* 4 days */
         %do j=1 %to 19; /* 19 scenes */
            var S&j._&k=[0,1]; /* Indicates if scene j is shot on day k.   */
         %end;
         %do i=1 %to 11; /* 11 actors */
            var A&i._&k=[0,1]; /* Indicates if actor i is present on day k.*/
         %end;
      %end;
```

```
      /* Every scene is shot exactly once.*/
      %do j=1 %to 19;
         gcc (
            %do k=1 %to 4;
               S&j._&k
            %end;
            )=( (1, 1, 1) );
      %end;

      /* At least 4 and at most 5 scenes are shot per day. */
      %do k=1 %to 4;
         gcc (
            %do j=1 %to 19;
               S&j._&k
            %end;
            )=( (1, 4, 5) );
      %end;

      /* Actors for a scene must be present on day of shooting.*/
      %do k=1 %to 4;
         %do j=1 %to 19;
            %do i=1 %to 11;
               %if (&&A&i._S&j>0) %then %do;
                  lincon S&j._&k <= A&i._&k;
               %end;
            %end;
         %end;
      %end;

      /* Symmetry breaking constraints. Without loss of any generality, we */
      /* can assume Scene1 to be shot on day 1, Scene2 to be shot on day 1 */
      /* or day 2, and Scene3 to be shot on either day 1, day 2 or day 3.  */
      lincon S1_1 = 1, S1_2 = 0, S1_3 = 0, S1_4 = 0,
             S2_3 = 0, S2_4 = 0, S3_4 = 0;

      /* If Scene2 is shot on day 1,                    */
      /* then Scene3 can be shot on day 1 or day 2. */
      lincon S2_1 + S3_3 <= 1;

      /* minimize total cost */
      lincon &obj >=
         %do i=1 %to 11; /* 11 actors */
            %do k=1 %to 4; /* 4 days */
               &&Cost&i*A&i._&k +
            %end;
         %end;
         0;
      run;

      /* when a solution is found,                        */
      /* &_ORCLP_ contains the string SOLUTIONS_FOUND   */
      %if %index(&_ORCLP_, SOLUTIONS_FOUND) %then %let clpreturn=SUCCESSFUL;
%mend sceneclp;
```

The %FINDMIN macro uses a bisection search method to compute the minimal cost. The %SCENE macro reads the SCENE data set and produces three sets of macro variables: Name$i$ is the name of the ith actor; Cost$i$ is the daily cost of the ith actor; and Ai_Sj $= 1$ if actor i appears in scene j, and 0 otherwise.

```
/* Bisection search method to determine the optimal objective value  */
%macro findmin(lb, ub);
   %do %while (&lb<&ub-1) ;
      %put Currently lb=&lb, ub=&ub..;
      %let newobj=%eval((&lb+&ub)/2);
      %let clpreturn=NOTFOUND;
      %sceneclp(obj=&newobj);
      %if &clpreturn=SUCCESSFUL %then %let ub=&newobj;
      %else %let lb=&newobj;
   %end;
   %sceneclp(obj=&ub);
   %put The minimum objective value within given range is &ub..;
   %put Any value less than &ub makes the problem infeasible. ;
%mend findmin;

%macro scene;
   /* Ai_Sj=1 if actor i appears in scene j    */
   /* Ai_Sj=0 otherwise                        */
   /* Initialize to 0                          */
   %do i=1 %to 11; /* 11 actors */
      %do j=1 %to 19; /* 19 scenes */
         %let A&i._S&j=0;
      %end;
   %end;

   data _null_;
      set scene;
      call symput("Name"||strip(_n_), Actor); /* read actor name */
      call symput("Cost"||strip(_n_), DailyFee); /* read actor cost */
      /* read whether an actor appears in a scene */
      %do i=1 %to 9; /* 9 scene variables in the data set */
         if S_Var&i > 0 then
            call symput("A"||strip(_n_)||"_S"||strip(S_Var&i), 1);
      %end;
   run;

   /* Find the minimum objective value.                  */
   /* Lowerbound: every actor appears on one day.        */
   /* Upperbound: every actor appears on all four days.  */
   %findmin(137739, 550956);

   /* Generate the Gantt charts */
   %gantt_gen;

%mend scene;
%scene;
```
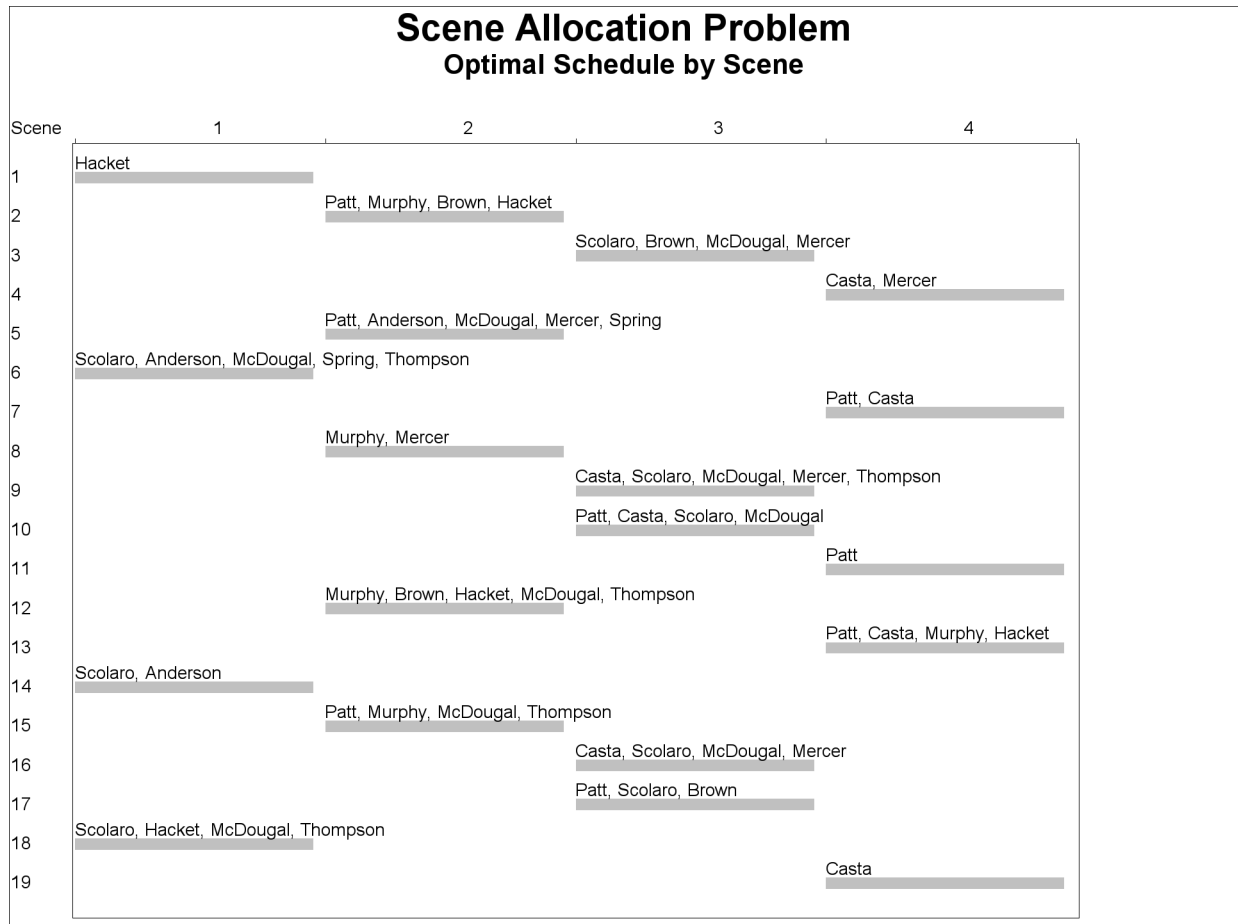
The optimal production cost is found to be 334,144. The corresponding actor schedules and scene schedules are displayed in and , respectively.

Output 3.6.2 Scene Allocation Problem: Actor Schedules

Output 3.6.3 Scene Allocation Problem: Scene Schedules

Scene Allocation Problem
Optimal Schedule by Scene

Scene	1	2	3	4
1	Hacket			
2		Patt, Murphy, Brown, Hacket		
3			Scolaro, Brown, McDougal, Mercer	
4				Casta, Mercer
5		Patt, Anderson, McDougal, Mercer, Spring		
6	Scolaro, Anderson, McDougal, Spring, Thompson			
7				Patt, Casta
8		Murphy, Mercer		
9			Casta, Scolaro, McDougal, Mercer, Thompson	
10			Patt, Casta, Scolaro, McDougal	
11				Patt
12		Murphy, Brown, Hacket, McDougal, Thompson		
13				Patt, Casta, Murphy, Hacket
14	Scolaro, Anderson			
15		Patt, Murphy, McDougal, Thompson		
16			Casta, Scolaro, McDougal, Mercer	
17			Patt, Scolaro, Brown	
18	Scolaro, Hacket, McDougal, Thompson			
19				Casta

Example 3.7: Car Sequencing Problem

This problem is an instance of a category of problems known as the car sequencing problem. There is a considerable amount of literature related to this problem (Dincbas, Simonis, and Van Hentenryck 1988; Gravel, Gagne, and Price 2005; Solnon et al. 2008).

A number of cars are to be produced on an assembly line where each car is customized with a specific set of options such as air-conditioning, sunroof, navigation, and so on. The assembly line moves through several workstations for installation of these options. The cars cannot be positioned randomly since each of these workstations have limited capacity and need time to set up the options as the assembly line is moving in front of the station. These capacity constraints are formalized using constraints of the form m out of N, which indicates that the workstation can install the option on m out of every sequence of N cars. The car sequencing problem is to determine a sequencing of the cars on the assembly line that satisfies the demand constraints for each set of car options and the capacity constraints for each workstation.

This example comes from Dincbas, Simonis, and Van Hentenryck (1988). Ten cars need to be customized with five possible options. A class of car is defined by a specific set of options; there are six classes of cars.

The instance data are presented in Table 3.9.

Table 3.9 The Instance Data

Option Name	Type	Capacity m/N	Car Class 1	2	3	4	5	6
Option 1	1	1/2	1	0	0	0	1	1
Option 2	2	2/3	0	0	1	1	0	1
Option 3	3	1/3	1	0	0	0	1	0
Option 4	4	2/5	1	1	0	1	0	0
Option 5	5	1/5	0	0	1	0	0	0
Number of Cars			**1**	**1**	**2**	**2**	**2**	**2**

For example, car class 4 requires installation of option 2 and option 4, and two cars of this class are required. The workstation for option 2 can process only two out of every sequence of three cars. The workstation for option 4 has even less capacity—two out of every five cars.

The instance data for this problem is used to create a SAS data set, which in turn is processed to generate the SAS macro variables shown in Table 3.10 that are used in the CLP procedure. The assembly line is treated as a sequence of slots, and each car must be allocated to a single slot.

Table 3.10 SAS Macro Variables

Macro Variable	Description	Value
Ncars	Number of cars (slots)	10
Nops	Number of options	5
Nclss	Number of classes	6
Max_1–Max_5	For each option, the maximum number of cars with that option in a block	1 2 1 2 1
Blsz_1–Blsz_5	For each option, the block size to which the maximum number refers	2 3 3 5 5
class_1–class_6	Index number of each class	1 2 3 4 5 6
cars_cls_1–cars_cls_6	Number of cars in each class	1 1 2 2 2 2
list_1–list_5	Class indicator list for each option; for example, classes 1, 5, and 6 that require option 1 (*list_1*)	list_1=1,0,0,0,1,1 list_2=0,0,1,1,0,1 list_3=1,0,0,0,1,0 list_4=1,1,0,1,0,0 list_5=0,0,1,0,0,0
cars_opts_1–cars_opts_5	Number of cars for each option (cars_opts_1 represents the number of cars that require option 1)	5 6 3 4 2

The decision variables for this problem are shown in Table 3.11.

Table 3.11 The Decision Variables

Variable Definition	Description
S_1–S_10=[1,6]	S_i is the class of cars assigned to slot i.
O_1_1–O_1_5=[0,1] ... O_10_1–O_10_5=[0,1]	O_i_j=1 if the class assigned to slot i needs option j.

In the following SAS statments, the workstation capacity constraints are expressed using a set of linear constraints for each workstation. The demand constraints for each car class are expressed using a single GCC constraint. The relationships between slot variables and option variables are expressed using an element constraint for each option variable. Finally, a set of redundant constraints are introduced to improve the efficiency of propagation. The idea behind the redundant constraint is the following realization: if the workstation for option j has capacity r out of s, then at most r cars in the sequence $(n - s + 1), \ldots, n$ can have option j where n is the total number of cars. Consequently at least $n_j - r$ cars in the sequence $1, \ldots, n - s$ must have option j, where n_j is the number of cars with option j. Generalizing this further, at least $n_j - k \times r$ cars in the sequence $1, \ldots, (n - k \times s)$ must have option j, $k = 1, \ldots, \lfloor n/s \rfloor$.

```
%macro car_sequencing(outdata);

proc clp out=&outdata varselect=minrmaxc findall;

   /* Declare Variables */
   var
      /* Slot variables: Si - class of car assigned to Slot i */
      %do i = 1 %to &Ncars;
          S_&i = [1, &Nclss]
      %end;

      /* Option variables: Oij
         - indicates if class assigned to Sloti needs Option j */
      %do i = 1 %to &Ncars;
         %do j = 1 %to &Nops;
             O_&i._&j = [0, 1]
         %end;
      %end;
      ;

   /* Capacity Constraints: for each option j */
   /* Installed in at most Max_j cars out of every sequence of BlSz_j cars */
   %do j = 1 %to &Nops;
      %do i = 0 %to %eval(&Ncars-&&BlSz_&j);
         lincon 0
         %do k=1 %to &&BlSz_&j;
            + O_%eval(&i+&k)_&j
         %end;
         <=&&Max_&j;
      %end;
   %end;
```

```
   /* Demand Constraints: for each class i */
   /* Exactly cars_cls_i cars */
   gcc (S_1-S_&Ncars) = (
      %do i = 1 %to &Nclss;
         (&&class_&i, &&cars_cls_&i, &&cars_cls_&i)
      %end;
      );

   /* Element Constraints: For each slot i and each option j */
   /* relate the slot variable to the option variables.     */
   /*  O_i_j is the S_i th element in list_j.                */
   %do i = 1 %to &Ncars;
      %do j = 1 %to &Nops;
         element (S_&i, (&&list_&j), O_&i._&j);
      %end;
   %end;


   /* Redundant Constraints to improve efficiency - for every  */
   /*    option j.                                             */
   /* At most &&Max_&j out of every sequence of &&BlSz_&j cars */
   /*    requires option j.                                    */
   /* All the other slots contain at least cars_opt_j - Max_j  */
   /*    cars with option j                                    */
   %do j = 1 %to &Nops;
      %do i = 1 %to %eval(&Ncars/&&BlSz_&j);
         lincon 0
         %do k=1 %to %eval(&Ncars-&i*&&BlSz_&j);
            + O_&k._&j
         %end;
         >= %eval(&&cars_opts_&j-&i*&&Max_&j);
      %end;
   %end;
run;

%mend;
%car_sequencing(sequence_out);
```

This problem has six solutions, as shown in

Output 3.7.1 Car Sequencing

Car Sequencing Problem
Assembly Sequence for All Six Solutions

Example 3.8: Round-Robin Problem

Round-robin tournaments (and variations of them) are a popular format in the scheduling of many sports league tournaments. In a single round-robin tournament, each team plays every other team just once. In a double round-robin (DRR) tournament, each team plays every other team twice: once at home and once away.

This particular example deals with a single round-robin tournament by modeling it as a scheduling CSP. A special case of a double round-robin tournament can be found in Example 3.12, "Scheduling a Major Basketball Conference" and features a different modeling approach.

Consider 14 teams that participate in a single round-robin tournament. Four rooms are provided for the tournament. Thus, $\binom{14}{2} = 91$ games and $\lceil \frac{91}{4} \rceil = 23$ time slots (rounds) need to be scheduled. Since each game requires two teams, a room, and an available time slot, you can regard each game as an activity, the two teams and the room as resources required by the activity, and the time slot as defined by the start and finish times of the activity.

In other words, you can treat this as a scheduling CSP with activities ACT_i_j where $i < j$, and resources TEAM1 through TEAM14 and ROOM1 through ROOM4. For a given i and j, activity ACT_i_j requires the resources TEAMi, TEAMj, and one of ROOM1 through ROOM4. The

resulting start time for activity A_i_j is the time slot for the game between TEAMi and TEAMj and the assigned ROOM is the venue for the game.

The following SAS macro, %ROUND_ROBIN, uses the CLP procedure to solve this problem. The %ROUND_ROBIN macro uses the number of teams as a parameter.

The ACTDATA= data set defines all activities ACT$\_i\_j$ with duration one. The RESOURCE statement declares the TEAM and ROOM resources. The REQUIRES statement defines the resource requirements for each activity ACT$\_i\_j$. The SCHEDULE statement defines the activity selection strategy as MINLS, which selects an activity with minimum late start time from the set of activities that begin prior to the earliest early finish time.

```
%macro round_robin(nteams);

   %let nrounds = %eval(%sysfunc(ceil((&nteams * (&nteams - 1)/2)/4)));

   data actdata;
      do i = 1 to &nteams - 1;
         do j = i + 1 to &nteams;
            _activity_ = compress('ACT_'||put(i,best.)||'_'||put(j,best.));
            _duration_ = 1;
            output;
         end;
      end;
   run;

   proc clp actdata = actdata schedule = schedule;
      schedule dur = &nrounds actselect=minls;
      resource (TEAM1-TEAM&nteams);
      resource (ROOM1-ROOM4);
      requires
         %do i = 1 %to &nteams - 1;
            %do j = &i + 1 %to &nteams;
               ACT_&i._&j = ( TEAM&i )
               ACT_&i._&j = ( TEAM&j )
               ACT_&i._&j = ( ROOM1, ROOM2, ROOM3, ROOM4)
            %end;
         %end;
      ;
      run;

   proc sort data=schedule;
      by start finish;
   run;

   %mend round_robin;

%round_robin(14);
```
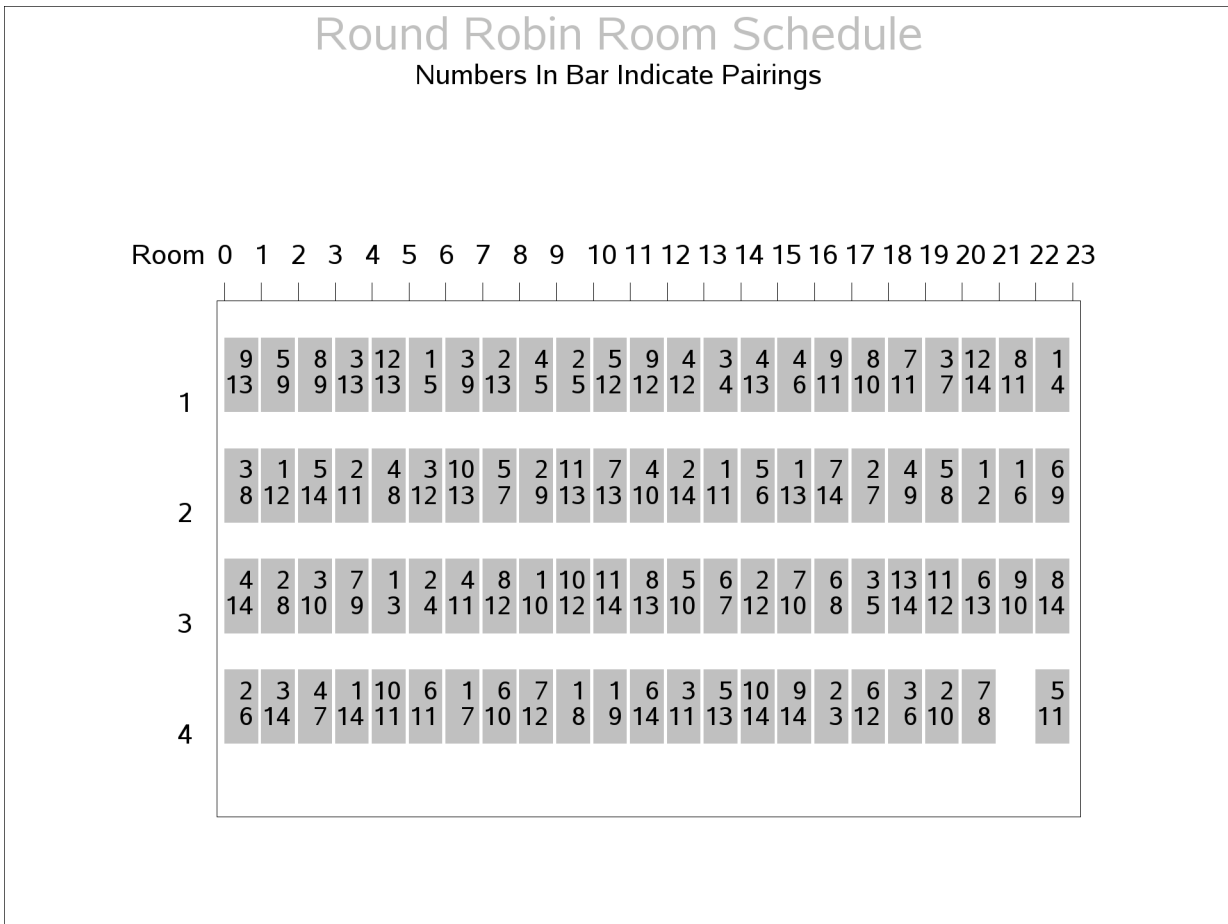
The resulting team schedule is displayed in Output 3.8.1. The vertical axis lists the teams, and the horizontal axis indicates the time slot of each game. The color of the bar indicates the room the game is played in, while the text above each bar identifies the opponent.

Output 3.8.1 Round Robin Team Schedule

Another view of the complete schedule is the room schedule, which is shown in Output 3.8.2. The vertical axis lists each room, and the horizontal axis indicates the time slot of each game. The numbers inside each bar identify the team pairings for the corresponding room and time slot.

Output 3.8.2 Round Robin Room Schedule

Round Robin Room Schedule
Numbers In Bar Indicate Pairings

Room	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23
1	9/13	5/9	8/9	3/13	12/13	1/5	3/9	2/13	4/5	2/5	5/12	9/12	4/12	3/4	4/13	4/6	9/11	8/10	7/11	3/7	12/14	8/11	1/4	
2	3/8	1/12	5/14	2/11	4/8	3/12	10/13	5/7	2/9	11/13	7/13	4/10	2/14	1/11	5/6	1/13	7/14	2/7	4/9	5/8	1/2	1/6	6/9	
3	4/14	2/8	3/10	7/9	1/3	2/4	4/11	8/12	1/10	10/12	11/14	8/13	5/10	6/7	2/12	7/10	6/8	3/5	13/14	11/12	6/13	9/10	8/14	
4	2/6	3/14	4/7	1/14	10/11	6/11	1/7	6/10	7/12	1/8	1/9	6/14	3/11	5/13	10/14	9/14	2/3	6/12	3/6	2/10	7/8			5/11

Example 3.9: Resource-Constrained Scheduling with Nonstandard Temporal Constraints

This example illustrates a real-life scheduling problem and is used as a benchmark problem in the constraint programming community. The problem is to schedule the construction of a five-segment bridge. (See Output 3.9.1.) It comes from a Ph.D. dissertation on scheduling problems (Bartusch 1983).

Output 3.9.1 The Bridge Problem

The project consists of 44 tasks and a set of constraints that relate these tasks. Table 3.12 displays the activity information, standard precedence constraints, and resource constraints. The sharing of a unary resource by multiple activities results in the resource constraints being disjunctive in nature.

Table 3.12 Data for Bridge Construction

Activity	Description	Duration	Predecessors	Resource
pa	beginning of project	0		
a1	excavation (abutment 1)	4	pa	excavator
a2	excavation (pillar 1)	2	pa	excavator
a3	excavation (pillar 2)	2	pa	excavator
a4	excavation (pillar 3)	2	pa	excavator
a5	excavation (pillar 4)	2	pa	excavator
a6	excavation (abutment 2)	5	pa	excavator
p1	foundation piles 2	20	a3	pile driver
p2	foundation piles 3	13	a4	pile driver
ue	erection of temporary housing	10	pa	
s1	formwork (abutment 1)	8	a1	carpentry
s2	formwork (pillar 1)	4	a2	carpentry

Table 3.12 *continued*

Activity	Description	Duration	Predecessors	Resource
s3	formwork (pillar 2)	4	p1	carpentry
s4	formwork (pillar 3)	4	p2	carpentry
s5	formwork (pillar 4)	4	a5	carpentry
s6	formwork (abutment 2)	10	a6	carpentry
b1	concrete foundation (abutment 1)	1	s1	concrete mixer
b2	concrete foundation (pillar 1)	1	s2	concrete mixer
b3	concrete foundation (pillar 2)	1	s3	concrete mixer
b4	concrete foundation (pillar 3)	1	s4	concrete mixer
b5	concrete foundation (pillar 4)	1	s5	concrete mixer
b6	concrete foundation (abutment 2)	1	s6	concrete mixer
ab1	concrete setting time (abutment 1)	1	b1	
ab2	concrete setting time (pillar 1)	1	b2	
ab3	concrete setting time (pillar 2)	1	b3	
ab4	concrete setting time (pillar 3)	1	b4	
ab5	concrete setting time (pillar 4)	1	b5	
ab6	concrete setting time (abutment 2)	1	b6	
m1	masonry work (abutment 1)	16	ab1	bricklaying
m2	masonry work (pillar 1)	8	ab2	bricklaying
m3	masonry work (pillar 2)	8	ab3	bricklaying
m4	masonry work (pillar 3)	8	ab4	bricklaying
m5	masonry work (pillar 4)	8	ab5	bricklaying
m6	masonry work (abutment 2)	20	ab6	bricklaying
l	delivery of the preformed bearers	2		crane
t1	positioning (preformed bearer 1)	12	m1, m2, l	crane
t2	positioning (preformed bearer 2)	12	m2, m3, l	crane
t3	positioning (preformed bearer 3)	12	m3, m4, l	crane
t4	positioning (preformed bearer 4)	12	m4, m5, l	crane
t5	positioning (preformed bearer 5)	12	m5, m6, l	crane
ua	removal of the temporary housing	10		
v1	filling 1	15	t1	caterpillar
v2	filling 2	10	t5	caterpillar
pe	end of project	0	t2, t3, t4, v1, v2, ua	

Output 3.9.2 shows a network diagram that illustrates the precedence constraints in this problem. Each node represents an activity and gives the activity code, truncated description, duration, and the required resource, if any. The network diagram is generated using the SAS/OR NETDRAW procedure.

Output 3.9.2   Network Diagram for the Bridge Construction Project
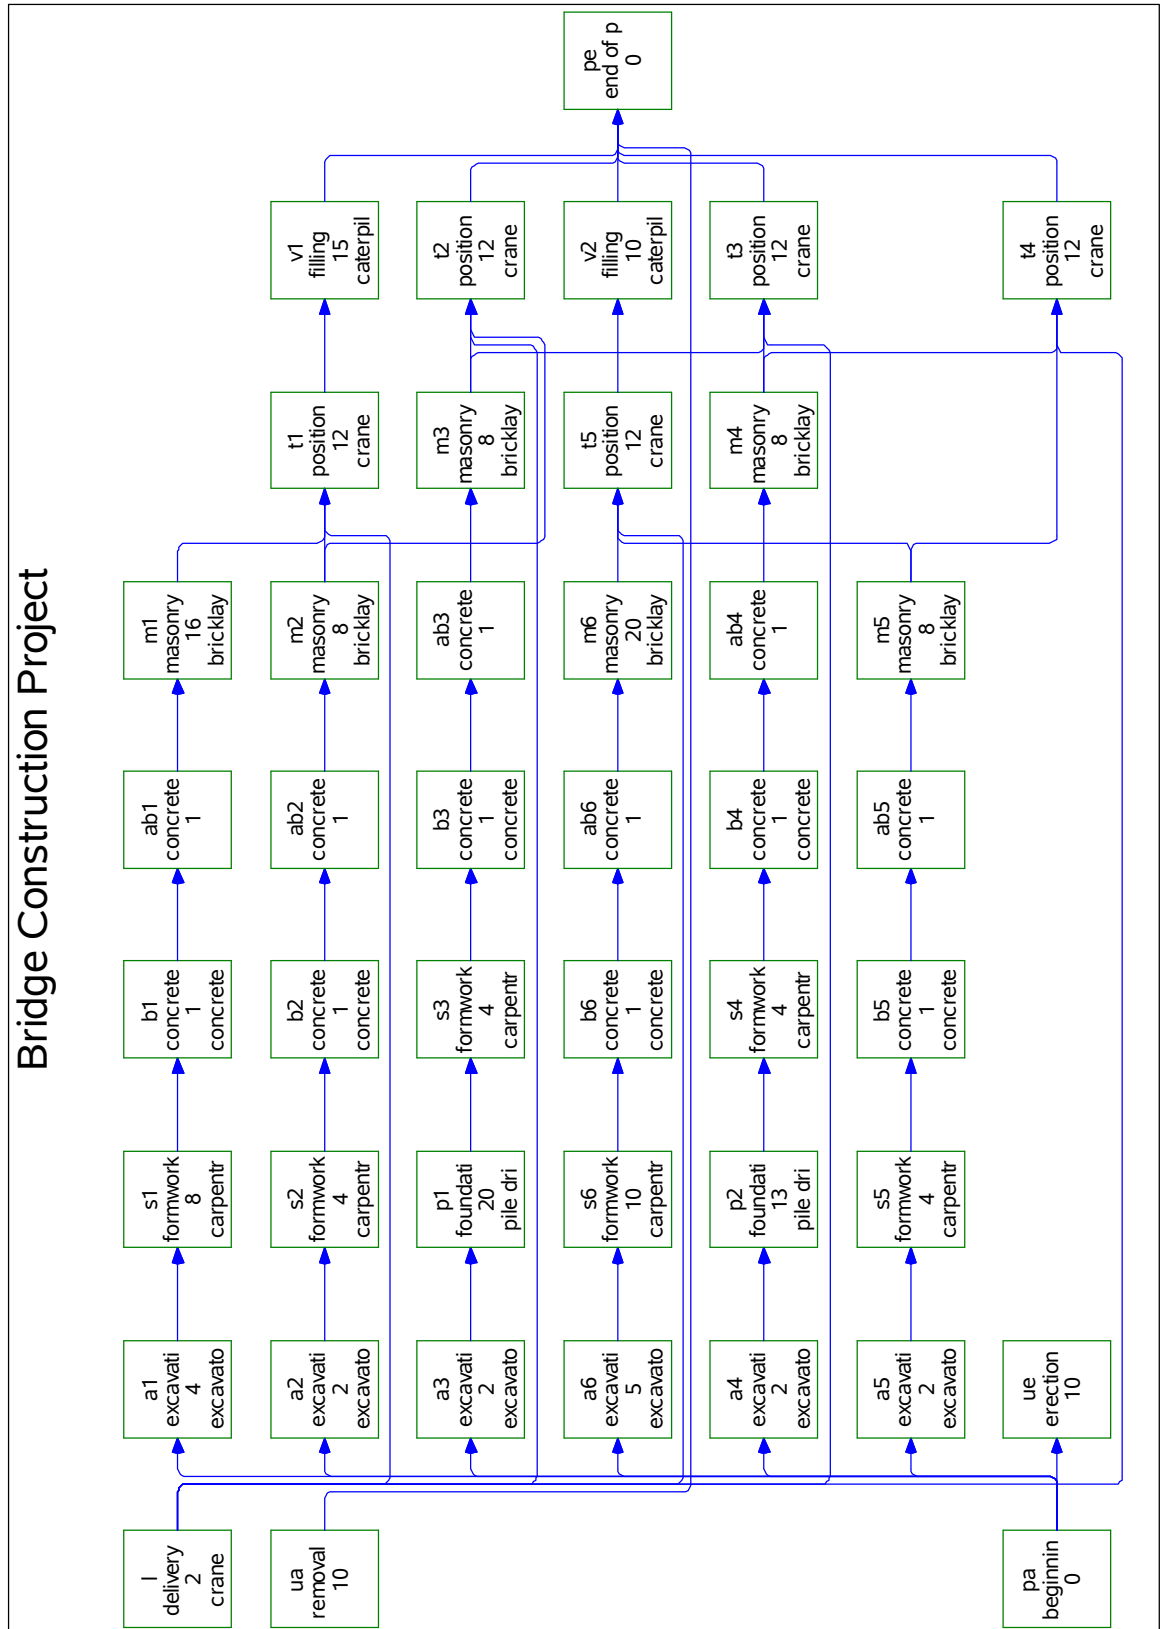
The following constraints are in addition to the standard precedence constraints:

1. The time between the completion of a particular formwork and the completion of its corresponding concrete foundation is at most four days:

$$f(si) \geq f(bi) - 4, \quad i = 1, \cdots, 6$$

2. There are at most three days between the end of a particular excavation (or foundation piles) and the beginning of the corresponding formwork:

$$f(ai) \geq s(si) - 3, \quad i = 1, 2, 5, 6$$

and

$$f(p1) \geq s(s3) - 3$$

$$f(p2) \geq s(s4) - 3$$

3. The formworks must start at least six days after the beginning of the erection of the temporary housing:

$$s(si) \geq s(ue) + 6, \quad i = 1, \cdots, 6$$

4. The removal of the temporary housing can start at most two days before the end of the last masonry work:

$$s(ua) \geq f(mi) - 2, \quad i = 1, \cdots, 6$$

5. The delivery of the preformed bearers occurs exactly 30 days after the beginning of the project:

$$s(l) = s(pa) + 30$$

The following DATA step defines the data set bridge, which encapsulates all of the precedence constraints and also indicates the resources that are required by each activity. Note the use of the reserved variables _ACTIVITY_, _SUCCESSOR_, _LAG_, and _LAGDUR_ to define the activity and precedence relationships. The list of reserved variables can be found in Table 3.5. The latter two variables are required for the nonstandard precedence constraints listed previously.

```
data bridge;
   format _ACTIVITY_ $3. _DESC_ $34. _RESOURCE_ $14.
          _SUCCESSOR_ $3. _LAG_ $3. ;
   input _ACTIVITY_ & _DESC_ & _DURATION_ _RESOURCE_ &
         _SUCCESSOR_ & _LAG_ & _LAGDUR_;
   _QTY_ = 1;
   datalines;
a1   excavation (abutment 1)               4   excavator       s1   .    .
a2   excavation (pillar 1)                 2   excavator       s2   .    .
a3   excavation (pillar 2)                 2   excavator       p1   .    .
a4   excavation (pillar 3)                 2   excavator       p2   .    .
a5   excavation (pillar 4)                 2   excavator       s5   .    .
a6   excavation (abutment 2)               5   excavator       s6   .    .
```

```
ab1   concrete setting time (abutment 1)  1  .                m1   .    .
ab2   concrete setting time (pillar 1)    1  .                m2   .    .
ab3   concrete setting time (pillar 2)    1  .                m3   .    .
ab4   concrete setting time (pillar 3)    1  .                m4   .    .
ab5   concrete setting time (pillar 4)    1  .                m5   .    .
ab6   concrete setting time (abutment 2)  1  .                m6   .    .
b1    concrete foundation (abutment 1)    1  concrete mixer   ab1  .    .
b1    concrete foundation (abutment 1)    1  concrete mixer   s1   ff   -4
b2    concrete foundation (pillar 1)      1  concrete mixer   ab2  .    .
b2    concrete foundation (pillar 1)      1  concrete mixer   s2   ff   -4
b3    concrete foundation (pillar 2)      1  concrete mixer   ab3  .    .
b3    concrete foundation (pillar 2)      1  concrete mixer   s3   ff   -4
b4    concrete foundation (pillar 3)      1  concrete mixer   ab4  .    .
b4    concrete foundation (pillar 3)      1  concrete mixer   s4   ff   -4
b5    concrete foundation (pillar 4)      1  concrete mixer   ab5  .    .
b5    concrete foundation (pillar 4)      1  concrete mixer   s5   ff   -4
b6    concrete foundation (abutment 2)    1  concrete mixer   ab6  .    .
b6    concrete foundation (abutment 2)    1  concrete mixer   s6   ff   -4
l     delivery of the preformed bearers   2  crane            t1   .    .
l     delivery of the preformed bearers   2  crane            t2   .    .
l     delivery of the preformed bearers   2  crane            t3   .    .
l     delivery of the preformed bearers   2  crane            t4   .    .
l     delivery of the preformed bearers   2  crane            t5   .    .
m1    masonry work (abutment 1)          16  bricklaying      t1   .    .
m1    masonry work (abutment 1)          16  bricklaying      ua   fs   -2
m2    masonry work (pillar 1)             8  bricklaying      t1   .    .
m2    masonry work (pillar 1)             8  bricklaying      t2   .    .
m2    masonry work (pillar 1)             8  bricklaying      ua   fs   -2
m3    masonry work (pillar 2)             8  bricklaying      t2   .    .
m3    masonry work (pillar 2)             8  bricklaying      t3   .    .
m3    masonry work (pillar 2)             8  bricklaying      ua   fs   -2
m4    masonry work (pillar 3)             8  bricklaying      t3   .    .
m4    masonry work (pillar 3)             8  bricklaying      t4   .    .
m4    masonry work (pillar 3)             8  bricklaying      ua   fs   -2
m5    masonry work (pillar 4)             8  bricklaying      t4   .    .
m5    masonry work (pillar 4)             8  bricklaying      t5   .    .
m5    masonry work (pillar 4)             8  bricklaying      ua   fs   -2
m6    masonry work (abutment 2)          20  bricklaying      t5   .    .
m6    masonry work (abutment 2)          20  bricklaying      ua   fs   -2
p1    foundation piles 2                 20  pile driver      s3   .    .
p2    foundation piles 3                 13  pile driver      s4   .    .
pa    beginning of project                0  .                a1   .    .
pa    beginning of project                0  .                a2   .    .
pa    beginning of project                0  .                a3   .    .
pa    beginning of project                0  .                a4   .    .
pa    beginning of project                0  .                a5   .    .
pa    beginning of project                0  .                a6   .    .
pa    beginning of project                0  .                l    fse  30
pa    beginning of project                0  .                ue   .    .
pe    end of project                      0  .                .    .    .
s1    formwork (abutment 1)               8  carpentry        b1   .    .
s1    formwork (abutment 1)               8  carpentry        a1   sf   -3
s2    formwork (pillar 1)                 4  carpentry        b2   .    .
s2    formwork (pillar 1)                 4  carpentry        a2   sf   -3
```

```
s3   formwork (pillar 2)              4  carpentry   b3   .    .
s3   formwork (pillar 2)              4  carpentry   p1   sf   -3
s4   formwork (pillar 3)              4  carpentry   b4   .    .
s4   formwork (pillar 3)              4  carpentry   p2   sf   -3
s5   formwork (pillar 4)              4  carpentry   b5   .    .
s5   formwork (pillar 4)              4  carpentry   a5   sf   -3
s6   formwork (abutment 2)           10  carpentry   b6   .    .
s6   formwork (abutment 2)           10  carpentry   a6   sf   -3
t1   positioning (preformed bearer 1) 12 crane       v1   .    .
t2   positioning (preformed bearer 2) 12 crane       pe   .    .
t3   positioning (preformed bearer 3) 12 crane       pe   .    .
t4   positioning (preformed bearer 4) 12 crane       pe   .    .
t5   positioning (preformed bearer 5) 12 crane       v2   .    .
ua   removal of the temporary housing 10 .           pe   .    .
ue   erection of temporary housing   10  .           .    .    .
ue   erection of temporary housing   10  .           s1   ss   6
ue   erection of temporary housing   10  .           s2   ss   6
ue   erection of temporary housing   10  .           s3   ss   6
ue   erection of temporary housing   10  .           s4   ss   6
ue   erection of temporary housing   10  .           s5   ss   6
ue   erection of temporary housing   10  .           s6   ss   6
v1   filling 1                       15  caterpillar pe   .    .
v2   filling 2                       10  caterpillar pe   .    .
;
```

The CLP procedure is then invoked by using the following statements with the SCHEDTIME= option.

```
/* invoke PROC CLP */
proc clp actdata=bridge schedtime=schedtime_bridge;
   schedule finish=104;
run;
```

The FINISH= option is specified to find a solution in 104 days, which also happens to be the optimal makespan.

The schedtime_bridge data set contains the activity start and finish times as computed by the CLP procedure. Since an activity gets assigned to at most one resource, it is possible to represent the complete schedule information more concisely by merging the schedtime_bridge data set with the bridge_info data set, as shown in the following statements.

```
/* Create Consolidated Schedule */
proc sql;
   create table bridge_info as
          select distinct _ACTIVITY_ as ACTIVITY format $3. length 3,
          _DESC_ as DESCRIPTION, _RESOURCE_ as RESOURCE from bridge;

proc sort data=schedtime_bridge;
   by ACTIVITY;
run;

data schedtime_bridge;
   merge bridge_info schedtime_bridge;
   by ACTIVITY;
run;
```

```
proc sort data=schedtime_bridge;
   by START FINISH;
run;

proc print data=schedtime_bridge noobs width=min;;
   title 'Bridge Construction Schedule';
run;
```

Output 3.9.3 shows the resulting merged data set.

Output 3.9.3 Bridge Construction Schedule

```
                         Bridge Construction Schedule

            D
            E
            S
     A      C                          R          S    D
     C      R                          E          O    U
     T      I                          S          L    R          F
     I      P                          O          U    A    S     I
     V      T                          U          T    T    T     N
O    I      I                          R          I    I    A     I
b    T      O                          C          O    O    R     S
s    Y      N                          E          N    N    T     H

 1   pa     beginning of project                          1    0    0    0
 2   a4     excavation (pillar 3)        excavator        1    2    0    2
 3   ue     erection of temporary housing                 1   10    0   10
 4   a5     excavation (pillar 4)        excavator        1    2    2    4
 5   p2     foundation piles 3           pile driver      1   13    2   15
 6   a2     excavation (pillar 1)        excavator        1    2    5    7
 7   s5     formwork (pillar 4)          carpentry        1    4    6   10
 8   a3     excavation (pillar 2)        excavator        1    2    7    9
 9   b5     concrete foundation (pillar 4)  concrete mixer 1    1   10   11
10   s2     formwork (pillar 1)          carpentry        1    4   10   14
11   ab5    concrete setting time (pillar 4)              1    1   11   12
12   a1     excavation (abutment 1)      excavator        1    4   12   16
13   m5     masonry work (pillar 4)      bricklaying      1    8   12   20
14   b2     concrete foundation (pillar 1)  concrete mixer 1    1   14   15
15   ab2    concrete setting time (pillar 1)              1    1   15   16
16   s4     formwork (pillar 3)          carpentry        1    4   15   19
17   p1     foundation piles 2           pile driver      1   20   15   35
18   b4     concrete foundation (pillar 3)  concrete mixer 1    1   19   20
19   a6     excavation (abutment 2)      excavator        1    5   19   24
20   s1     formwork (abutment 1)        carpentry        1    8   19   27
21   ab4    concrete setting time (pillar 3)              1    1   20   21
22   m2     masonry work (pillar 1)      bricklaying      1    8   20   28
23   b1     concrete foundation (abutment 1)  concrete mixer 1    1   27   28
24   s6     formwork (abutment 2)        carpentry        1   10   27   37
25   ab1    concrete setting time (abutment 1)            1    1   28   29
26   m4     masonry work (pillar 3)      bricklaying      1    8   28   36
27   l      delivery of the preformed bearers  crane      1    2   30   32
28   t4     positioning (preformed bearer 4)  crane       1   12   36   48
29   m1     masonry work (abutment 1)    bricklaying      1   16   36   52
```

Output 3.9.3 *continued*

```
                        Bridge Construction Schedule

            D
            E
            S
    A       C                                   R               S    D
    C       R                                   E               O    U
    T       I                                   S               L    R              F
    I       P                                   O               U    A    S         I
    V       T                                   U               T    T    T         N
O   I       I                                   R               I    I    A         I
b   T       O                                   C               O    O    R         S
s   Y       N                                   E               N    N    T         H

30  b6    concrete foundation (abutment 2)   concrete mixer  1    1    37   38
31  s3    formwork (pillar 2)                carpentry       1    4    37   41
32  ab6   concrete setting time (abutment 2)                 1    1    38   39
33  b3    concrete foundation (pillar 2)     concrete mixer  1    1    41   42
34  ab3   concrete setting time (pillar 2)                   1    1    42   43
35  m3    masonry work (pillar 2)            bricklaying     1    8    52   60
36  t1    positioning (preformed bearer 1)   crane           1   12    52   64
37  m6    masonry work (abutment 2)          bricklaying     1   20    60   80
38  t2    positioning (preformed bearer 2)   crane           1   12    64   76
39  v1    filling 1                          caterpillar     1   15    64   79
40  ua    removal of the temporary housing                   1   10    78   88
41  t5    positioning (preformed bearer 5)   crane           1   12    80   92
42  v2    filling 2                          caterpillar     1   10    92  102
43  t3    positioning (preformed bearer 3)   crane           1   12    92  104
44  pe    end of project                                     1    0   104  104
```

A Gantt chart of the schedule in Output 3.9.3, produced using the SAS/OR GANTT procedure, is displayed in Output 3.9.4. Each activity bar is color coded according to the resource associated with it. The legend identifies the name of the resource that is associated with each color.

Output 3.9.4 Gantt Chart for the Bridge Construction Project

PROC CLP - Solution to Bridge Problem:104 days (optimal)
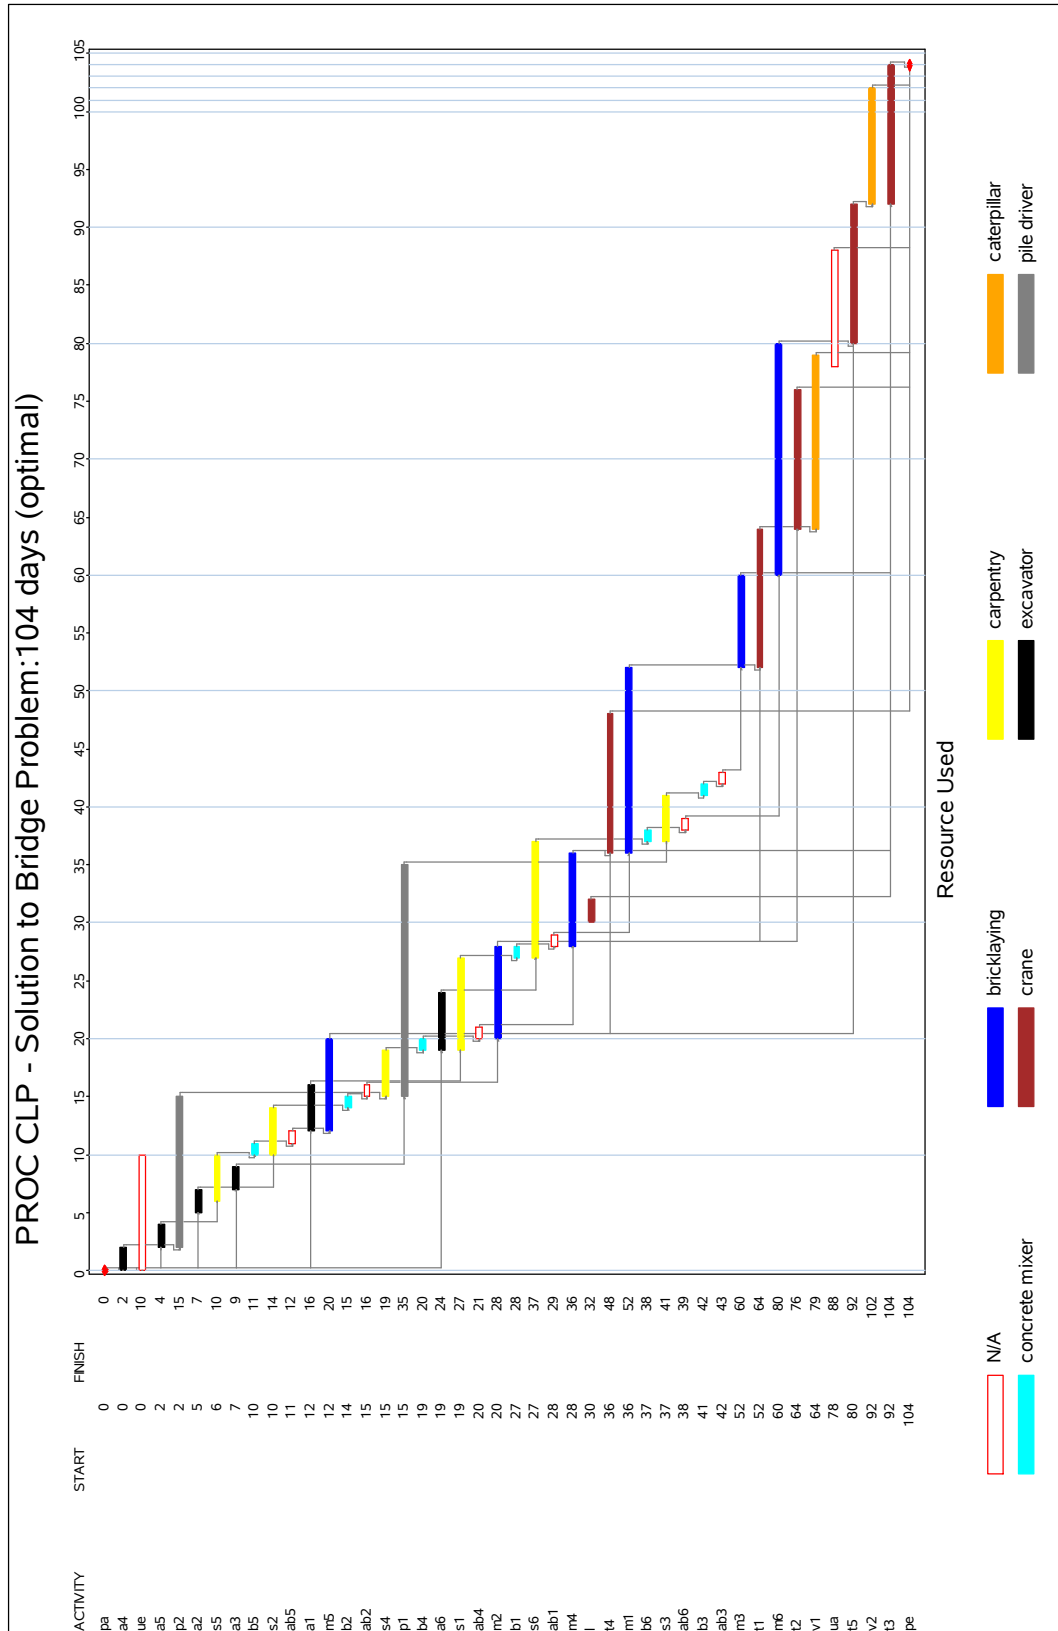
Example 3.10: Scheduling with Alternate Resources

This example shows an interesting job shop scheduling problem that illustrates the use of alternative resources. There are 90 jobs (J1–J90), each taking either one or two days, that need to be processed on one of ten machines (M0–M9). Not every machine can process every job. In addition, certain jobs also require one of seven operators (OP0–OP6). As with the machines, not every operator can be assigned to every job. There are no explicit precedence relationships in this example.

The machine and operator requirements for each job are shown in Output 3.10.1. Each row in the graph defines a resource requirement for up to three jobs that are identified in the columns Job1–Job3 to the left of the chart. The horizontal axis of the chart represents the resources and is split into two regions by a vertical line. The resources to the left of the divider are the machines, Mach0–Mach9, and the resources to the right of the divider are the operators, Oper0–Oper6. For each row on the chart, a bar on the chart represents a potential requirement for the corresponding resource listed above.

Each of the jobs listed in columns Job1–Job3 can be processed on one of the machines in Mach0–Mach9 and requires the assistance of one of the operators in Oper0–Oper6 while being processed. An eligible resource is represented by a bar, and the length of the bar indicates the duration of the job.

For example, row five specifies that job number 7 can be processed on machine 6, 7, or 8 and additionally requires either operator 5 or operator 6 in order to be processed. The next row indicates that jobs 8 and 9 can also be processed on the same set of machines. However, they do not require any operator assistance.

Output 3.10.1 Machine and Operator Requirements

Job1	Job2	Job3	Mach0	Mach1	Mach2	Mach3	Mach4	Mach5	Mach6	Mach7	Mach8	Mach9	Oper0	Oper1	Oper2	Oper3	Oper4	Oper5	Oper6
1			□	□									□	□	□				
2			□	□	□								□	□	□				
3				□	□	□							□	□	□				
4	5	6			□	□	□								□	□	□	□	
7								□	□	□								□	□
8	9							□	□	□									
10									□	□	□								
11				□	□								□	□	□	□			
12	13	14							□	□	□		□						
15	16						□	□									□	□	□
17					□	□											□	□	□
18					□	□													
19			□	□	□														
20			□	□															
21	22		□	□									□	□	□	□			
23					□									□	□	□			
24	25	36								□	□								
26	35	75						□	□										
27	34	74						□	□								□	□	
28						□	□										□	□	
29						□	□									□	□		
30						□										□	□		
31						□									□	□	□		
32							□	□							□	□			
33							□	□										□	□
37	76	77								□	□								
38	39		□	□									□	□					
40					□											□			
41								□	□										□
42	62	82						□	□										
43	44	63								□	□								
45	46	65	□	□													□	□	
47	67				□	□										□	□		
48	68				□	□										□			
49	50	69					□	□					□	□					
51	52					□	□						□	□					
53								□	□				□						
54							□	□											□
55	56	57							□	□	□								
58	59	78	□	□	□											□	□		
60	80			□	□													□	□
61	81							□	□									□	□
64	83	84								□	□								
66			□	□													□	□	
70							□						□	□					
71							□						□	□	□				
72	73					□	□						□	□	□				
79				□	□											□	□		
85			□													□			
86				□													□		
87					□													□	
88						□							□						
89							□									□			
90								□								□			
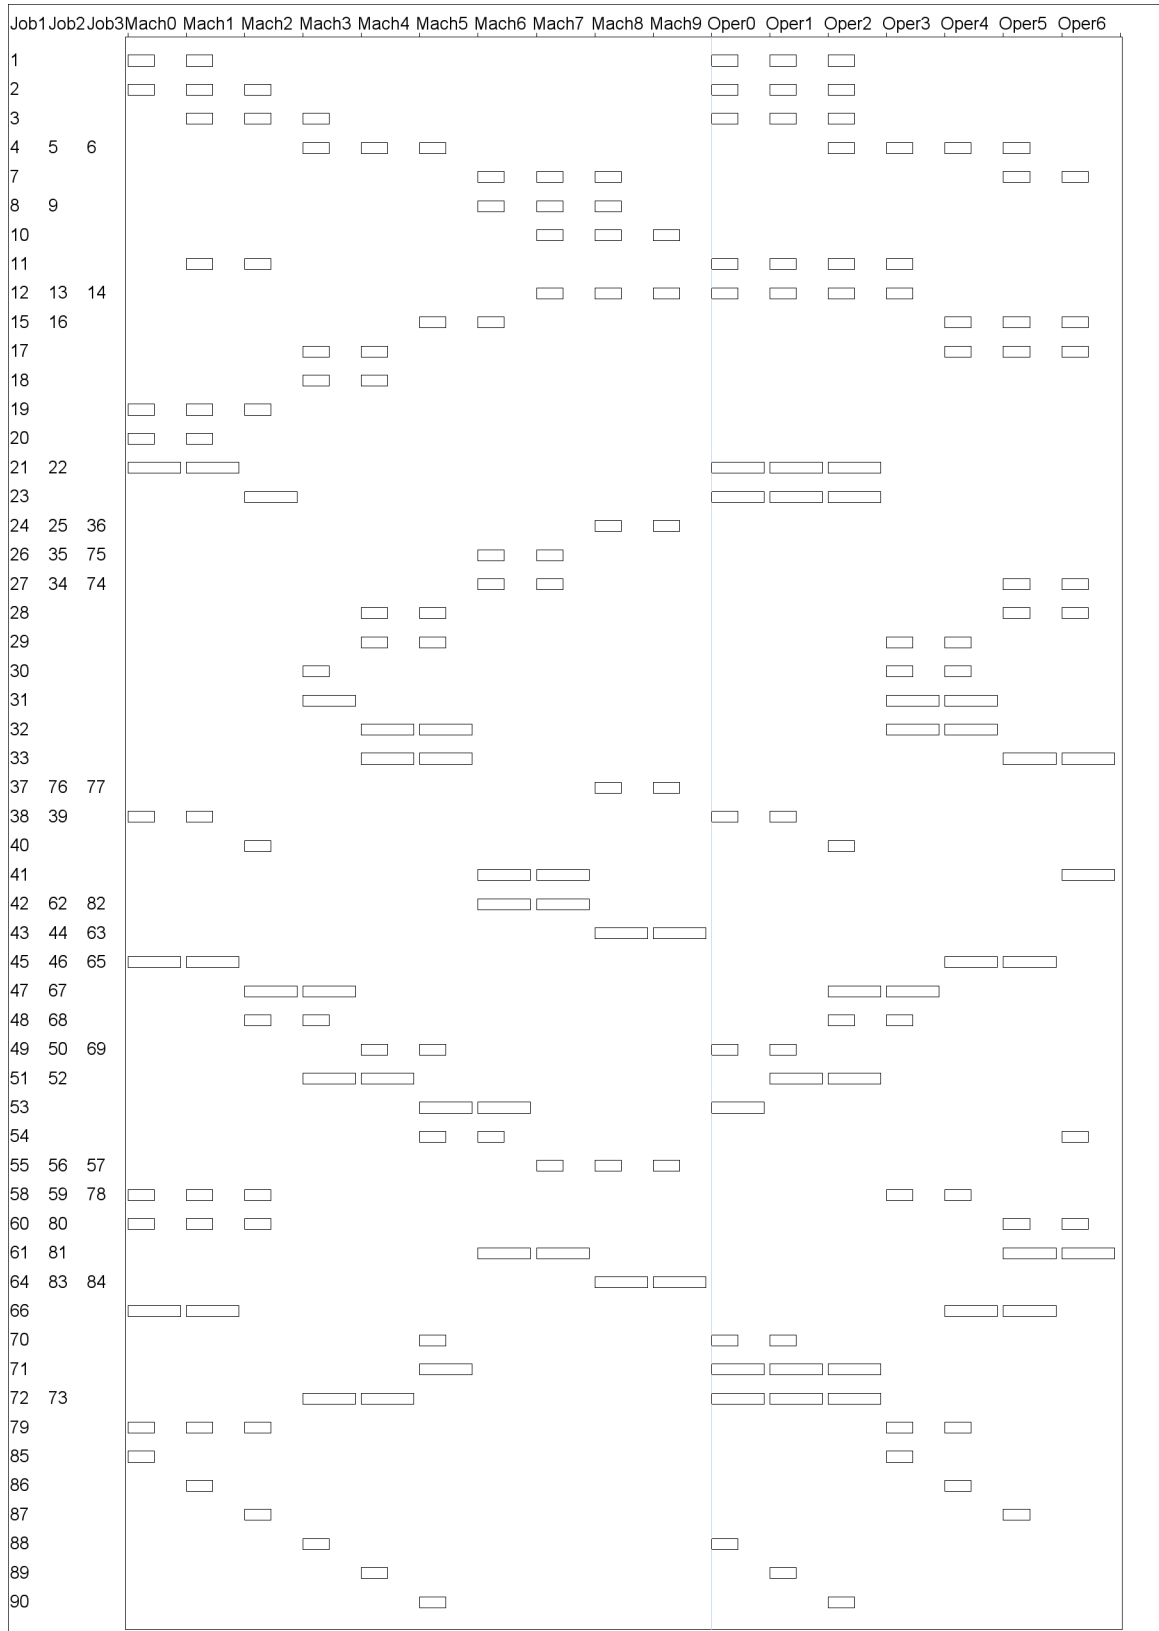
The CLP procedure is invoked by using the following statements with DUR=12 in the SCHEDULE statement to obtain a 12-day solution that is also known to be optimal. In order to obtain the optimal solution, it is necessary to invoke the edge-finding consistency routines, which are activated with the EDGEFINDER option in the SCHEDULE statement. The activity selection strategy is specified as DMINLS, which selects the activity with the earliest late start time. Activities with identical resource requirements are grouped together in the REQUIRES statement.

```
proc clp dom=[0,12] restarts=500 dpr=6 showprogress
   schedtime=schedtime_altres schedres=schedres_altres;
   schedule start=0 dur=12 actselect=dminls edgefinder;
   activity (J1-J20 J24-J30 J34-J40 J48-J50 J54-J60
            J68-J70 J74-J80 J85-J90) = (1)     /* one day jobs */
            (J21-J23 J31-J33 J41-J47 J51-J53 J61-J67
             J71-J73 J81-J84) = (2);           /* two day jobs */
   resource (M0-M9) (OP0-OP6);
   requires
      /* machine requirements */
      (J85) = (M0)
      (J1 J20 J21 J22 J38 J39 J45 J46 J65 J66) = (M0, M1)
      (J19 J2 J58 J59 J60 J78 J79 J80) = (M0, M1, M2)
      (J86) = (M1)
      (J11) = (M1, M2)
      (J3) = (M1, M2, M3)
      (J23 J40 J87) = (M2)
      (J47 J48 J67 J68) = (M2, M3)
      (J30 J31 J88) = (M3)
      (J17 J18 J51 J52 J72 J73) = (M3, M4)
      (J4 J5 J6) = (M3, M4, M5)
      (J89) = (M4)
      (J28 J29 J32 J33 J49 J50 J69) = (M4, M5)
      (J70 J71 J90) = (M5)
      (J15 J16 J53 J54) = (M5, M6)
      (J26 J27 J34 J35 J41 J42 J61 J62 J74 J75 J81 J82) = (M6, M7)
      (J7 J8 J9) = (M6, M7, M8)
      (J10 J12 J13 J14 J55 J56 J57) = (M7, M8, M9)
      (J24 J25 J36 J37 J43 J44 J63 J64 J76 J77 J83 J84) = (M8, M9)
      /* operator requirements */
      (J53 J88) = (OP0)
      (J38 J39 J49 J50 J69 J70) = (OP0, OP1)
      (J1 J2 J21 J22 J23 J3 J71 J72 J73) = (OP0, OP1, OP2)
      (J11 J12 J13 J14) = (OP0, OP1, OP2, OP3)
      (J89) = (OP1)
      (J51 J52) = (OP1, OP2)
      (J40 J90) = (OP2)
      (J47 J48 J67 J68) = (OP2, OP3)
      (J4 J5 J6) = (OP2, OP3, OP4, OP5)
      (J85) = (OP3)
      (J29 J30 J31 J32 J58 J59 J78 J79) = (OP3, OP4)
      (J86) = (OP4)
      (J45 J46 J65 J66) = (OP4, OP5)
      (J15 J16 J17) = (OP4, OP5, OP6)
      (J87) = (OP5)
      (J27 J28 J33 J34 J60 J61 J7 J74 J80 J81) = (OP5, OP6)
      (J41 J54) = (OP6);
   run;
```

The resulting schedule is shown in a series of Gantt charts that are displayed in Output 3.10.2 and Output 3.10.3. In each of these Gantt charts, the vertical axis lists the different jobs, the horizontal bar represents the start and finish times for each of the jobs, and the text above each bar identifies the machine that the job is being processed on. Output 3.10.2 displays the schedule for the operator-assisted tasks (one for each operator), while Output 3.10.3 shows the schedule for automated tasks (that is, those tasks that do not require operator intervention).

Output 3.10.2 Operator-Assisted Jobs Schedule

Schedule for Operator OP0
Machine Identified Above Bar

Schedule for Operator OP1
Machine Identified Above Bar

Output 3.10.2 *continued*

Schedule for Operator OP2
Machine Identified Above Bar

Schedule for Operator OP3
Machine Identified Above Bar

Schedule for Operator OP4
Machine Identified Above Bar

Output 3.10.2 *continued*

Schedule for Operator OP5
Machine Identified Above Bar

Schedule for Operator OP6
Machine Identified Above Bar

Output 3.10.3 Automated Jobs Schedule

A more interesting Gantt chart is that of the resource schedule by machine, as shown in Output 3.10.4. This chart displays the schedule for each machine. Every row corresponds to a machine. Every bar on each row consists of multiple segments, and every segment represents a job that is processed on the machine. Each segment is also coded according to the operator assigned to it. The mapping of the coding is indicated in the legend. It is evident that the schedule is optimal since none of the machines or operators are idle at any time during the schedule.

Output 3.10.4 Machine Schedule

Example 3.11: 10×10 Job Shop Scheduling Problem

This example is a job shop scheduling problem from Lawrence (1984). This test is also known as LA19 in the literature, and its optimal makespan is known to be 842 (Applegate and Cook 1991). There are 10 jobs (J1–J10) and 10 machines (M0–M9). Every job must be processed on each of the 10 machines in a predefined sequence. The objective is to minimize the completion time of the last job to be processed, known as the makespan. The jobs are described in the data set raw by using the following statements.

```
/* jobs specification */
data raw (drop=i mid);
   do i=1 to 10;
      input mid _DURATION_ @;
      _RESOURCE_=compress('M'||put(mid,best.));
      output;
   end;
   datalines;
2  44  3   5  5  58  4  97  0   9  7  84  8  77  9  96  1  58  6  89
4  15  7  31  1  87  8  57  0  77  3  85  2  81  5  39  9  73  6  21
9  82  6  22  4  10  3  70  1  49  0  40  8  34  2  48  7  80  5  71
1  91  2  17  7  62  5  75  8  47  4  11  3   7  6  72  9  35  0  55
6  71  1  90  3  75  0  64  2  94  8  15  4  12  7  67  9  20  5  50
7  70  5  93  8  77  2  29  4  58  6  93  3  68  1  57  9   7  0  52
6  87  1  63  4  26  5   6  2  82  3  27  7  56  8  48  9  36  0  95
0  36  5  15  8  41  9  78  3  76  6  84  4  30  7  76  2  36  1   8
5  88  2  81  3  13  6  82  4  54  7  13  8  29  9  40  1  78  0  75
9  88  4  54  6  64  7  32  0  52  2   6  8  54  5  82  3   6  1  26
;
```

Each row in the DATALINES section specifies a job by 10 pairs of consecutive numbers. Each pair of numbers defines one task of the job, which represents the processing of a job on a machine. For each pair, the first number identifies the machine it executes on, and the second number is the duration. The order of the 10 pairs defines the sequence of the tasks for a job.

The following statements create the Activity data set actdata, which defines the activities, durations, and precedence constraints:

```
/* create the Activity data set */
data actdata (drop= i j);
   format _ACTIVITY_ $8. _SUCCESSOR_ $8.;
   set raw;
   _QTY_ = 1;
   i=mod(_n_-1,10)+1;
   j=int((_n_-1)/10)+1;
   _ACTIVITY_ = compress('J'||put(j,best.)||'P'||put(i,best.));
   JOB=j;
   TASK=i;
   if i LT 10 then
      _SUCCESSOR_ = compress('J'||put(j,best.)||'P'||put((i+1),best.));
   else
      _SUCCESSOR_ = ' ';
   output;
run;
```

Had there been sufficient machine capacity, the jobs could have been processed according to a schedule as shown in Output 3.11.1. The minimum makespan would be 617—the time it takes to complete Job 1.

Output 3.11.1 Gantt Chart: Schedule for the Unconstrained Problem
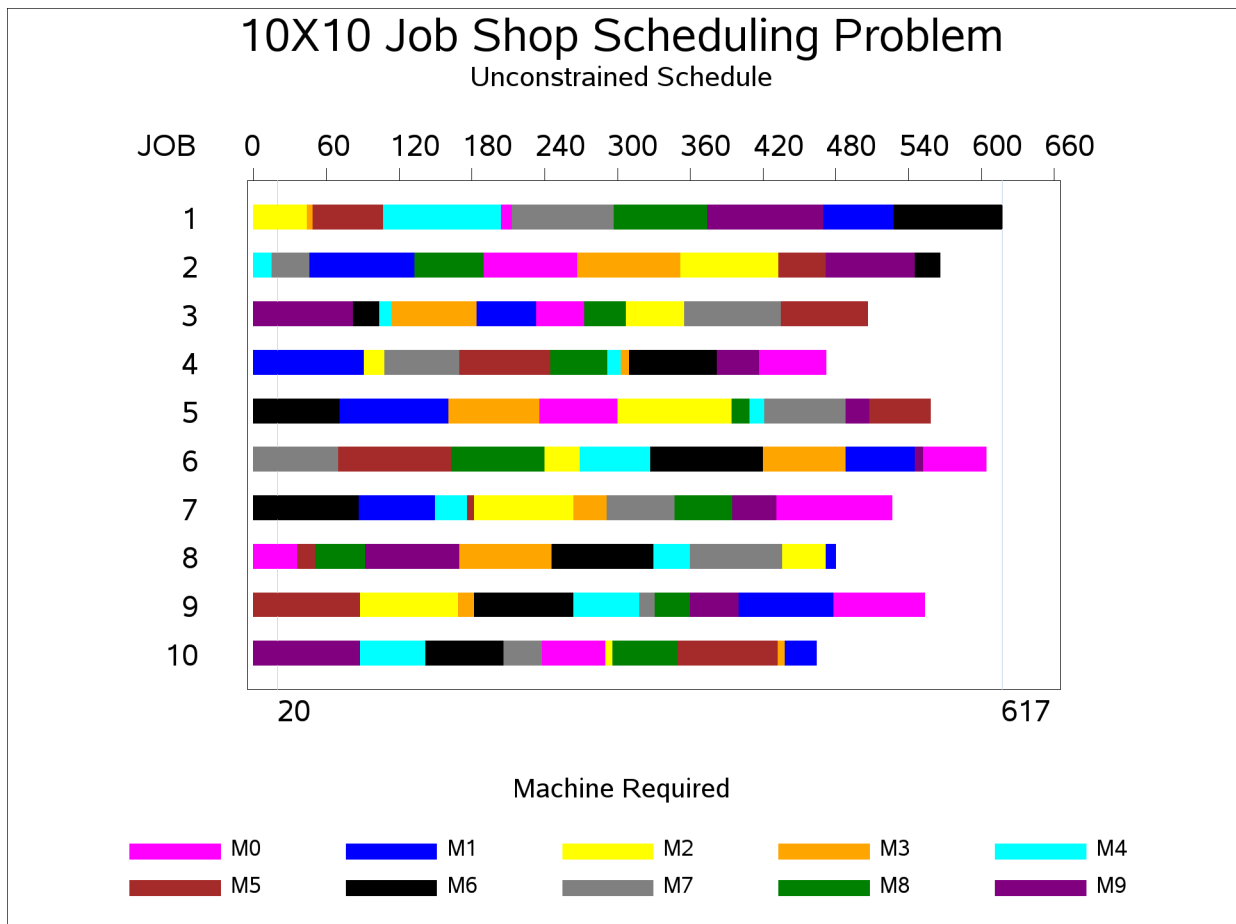
This schedule is infeasible when there is only a single instance of each machine. For example, at time period 20, the schedule requires two instances of each of the machines M6, M7, and M9.

In order to solve the resource-constrained schedule, the CLP procedure is invoked by using the following statements:

```
proc clp domain=[0,842]
         actdata=actdata
         schedout=sched_jobshop
         dpr=50
         restarts=150
         showprogress;
    schedule dur=842 edgefinder nf=1 nl=1;
run;
```

The edge-finder algorithm is activated with the EDGEFINDER option in the SCHEDULE statement. In addition, the edge-finding extensions for detecting whether a job cannot be the first or cannot be the last to be processed on a particular machine are invoked with the NF= and NL= options, respectively, in the SCHEDULE statement. The default activity selection and activity assignment strategies are used. A restart heuristic is used as the look-back method to handle recovery from failures. The DPR= option specifies that a total restart be performed after encountering 50 failures, and the RESTARTS= option limits the number of restarts to 150.

The resulting 842-time-period schedule is displayed in Output 3.11.2. Each row represents a job. Each segment represents a task (the processing of a job on a machine), which is also coded according to the executing machine. The mapping of the coding is indicated in the legend. Note that no machine is used by more than one job at any point in time.

Output 3.11.2 Gantt Chart: Optimal Resource-Constrained Schedule

Example 3.12: Scheduling a Major Basketball Conference

Example 1.8 illustrated how you could use the CLP procedure to solve a single round-robin problem by modeling it as a scheduling CSP. This example illustrates an alternate way of modeling and solving a well-known double round-robin problem using the CLP procedure. This example is based on the work of Nemhauser and Trick (1998) and deals with scheduling the Atlantic Coast Conference (ACC) Men's Basketball games for the 1997–1998 season.

A temporally dense double round-robin (DDRR) for n teams is a double round-robin in which the $n(n-1)$ games are played over a minimal number of dates or time slots. If n is even, the number of slots is $2(n-1)$ and each team plays in every time slot. If n is odd, the number of slots is $2n$ and $(n-1)$ teams play in each time slot. In the latter case, each time slot has a team with a bye, and each team has two byes for the season.

The Atlantic Coast Conference (ACC) 1997–1998 men's basketball scheduling problem as described in Nemhauser and Trick (1998) and Henz (2001) is a DDRR that consists of the following nine teams with their abbreviated team name and team number shown in parenthesis: Clemson (Clem 1)), Duke (Duke 2), Florida State (FSU 3), Georgia Tech (GT 4), Maryland (UMD 5), North Carolina (UNC 6) NC State (NCSU 7) Virginia (UVA 8), and Wake Forest (Wake 9).

The general objective is to schedule the DDRR to span the months of January and February and possibly include a game in December or March or both. In general, each team plays twice a week— typically Wednesday and Saturday. Although the actual day might differ, these two time slots are referred to as the "weekday slot" and the "weekend slot." Since there are an odd number of teams, there is a team with a bye in each slot and four games in each slot, resulting in a schedule that requires 18 time slots or nine weeks. The last time slot must be a weekend slot, which implies the first slot is a weekday slot. The first slot, denoted slot 1, corresponds to the last weekday slot of December 1997, and the final slot, slot 18, corresponds to the first weekend slot of March 1998. Each team plays eight home games and eight away games, and has two byes.

In addition there are several other constraints that must be satisfied. This example uses the following criteria employed by Nemhauser and Trick (1998) as presented by Henz (2001).

1. **Mirroring:** The dates are grouped into pairs $(r1, r2)$, such that each team gets to play against the same team in dates $r1$ and $r2$. Such a grouping is called a mirroring scheme. A separation of nine slots can be achieved by mirroring a round-robin schedule; while this separation is desirable, it is not possible for this problem.

   Nemhauser and Trick fix the mirroring scheme to

   $$m = (1, 8), (2, 9), (3, 12), (4, 13), (5, 14), (6, 15), (7, 16), (10, 17), (11, 18)$$

   in order to satisfy the constraints that UNC and Duke play in time slots 11 and 18. (See criterion 9.)

2. **Initial and final home and away games:** Every team must play at home on at least one of the first three dates. Every team must play at home on at least one of the last three dates. No team can play away on both of the last two dates.

3. **Home/away/bye pattern:** No team can have more than two away games in a row. No team can have more than two home games in a row. No team can have more than three away games or byes in a row. No team can have more than four home games or byes in a row.

4. **Weekend pattern:** Of the nine weekends, each team plays four at home, four away, and has one bye.

5. **First weekends:** Each team must have home games or byes on at least two of the first five weekends.

6. **Rival matches:** Every team except FSU has a traditional rival. The rival pairs are Clem-GT, Duke-UNC, UMD-UVA, and NCSU-Wake. In the last date, every team except FSU plays against its rival, unless it plays against FSU or has a bye.

7. **Popular matches in February:** The following pairings must occur at least once in dates 11 to 18: Duke-GT, Duke-Wake, GT-UNC, UNC-Wake.

8. **Opponent sequence:** No team plays in two consecutive away dates against Duke and UNC. No team plays in three consecutive dates against Duke, UNC, and Wake (independent of the home or away status).

9. **Idiosyncrasies:** UNC plays its rival Duke in the last date and in date 11. UNC plays Clem in the second date. Duke has a bye in date 16. Wake does not play home in date 17. Wake has a bye in the first date. Clem, Duke, UMD and Wake do not play away in the last date. Clem, FSU, and GT do not play away in the first date. Neither FSU nor NCSU has a bye in last date. UNC does not have a bye in the first date.

Previous work for solving round-robin problems, including that of Nemhauser and Trick (1998) and Henz (2001), have used a general three-phase framework for finding good schedules.

1. pattern generation

2. pattern set generation

3. timetable generation

A pattern is a valid sequence of home, away, and bye games for a given team for the entire season. For example, the following is a valid pattern:

A H B A H H A H A A H B H A A H H A

For this example, patterns that satisfy criterion 1 through criterion 5 and some constraints in criterion 9 are generated using the CLP procedure with the SAS macro %PATTERNS.

```
/**************************************************************/
/* First, find all possible patterns. Consider only time     */
/* constraints at this point. A pattern should be suitable    */
/* for any team. Do not consider individual teams yet.        */
/**************************************************************/
%macro patterns();

proc clp out=all_patterns findall;
   /* For date 1 to 18. */
   %do j = 1 %to 18;
      var h&j = [0, 1]; /* home */
      var a&j = [0, 1]; /* away */
      var b&j = [0, 1]; /* bye */

      /* A team is either home, away, or bye. */
      lincon h&j + a&j + b&j=1;
   %end;
```

```
/*--------------------------------------------------------------*/
/* Criterion 1 - Mirroring Scheme                               */
/*--------------------------------------------------------------*/
/* The dates are grouped into pairs (j, j1), such that each     */
/* team plays the same opponent on dates j and j1.              */
/* A home game on date j will be an away game on date j1         */
%do j = 1 %to 18;
   %do j1 = %eval(&j+1) %to 18;
      %if ( &j=1 and &j1=8 ) or ( &j=2 and &j1=9 ) or
          ( &j=3 and &j1=12 ) or ( &j=4 and &j1=13 ) or
          ( &j=5 and &j1=14 ) or ( &j=6 and &j1=15 ) or
          ( &j=7 and &j1=16 ) or ( &j=10 and &j1=17 ) or
          ( &j=11 and &j1=18 ) %then
         lincon h&j = a&j1, a&j = h&j1, b&j = b&j1;;
   %end;
%end;


/*--------------------------------------------------------------*/
/* Criterion 2 - Initial and Final Home and Away Games          */
/*--------------------------------------------------------------*/
/* Every team must play home on at least one of the first three dates. */
lincon h1 + h2 + h3 >= 1;

/* Every team must play home on at least one of the last three dates. */
lincon h16 + h17 + h18 >= 1;

/* No team can play away on both last two dates. */
lincon a17 + a18 < 2;


/*--------------------------------------------------------------*/
/* Criterion 3 - Home/Away/Bye Pattern                          */
/*--------------------------------------------------------------*/
%do j = 1 %to 16;
   /* No team can have more than two away matches in a row.*/
   lincon a&j + a%eval(&j+1) + a%eval(&j+2) < 3;
   /* No team can have more than two home matches in a row.*/
   lincon h&j + h%eval(&j+1) + h%eval(&j+2) < 3;
%end;

/* No team can have more than three away matches or byes in a row.*/
%do j = 1 %to 15;
   lincon a&j + b&j + a%eval(&j+1) + b%eval(&j+1) + a%eval(&j+2)
          + b%eval(&j+2) + a%eval(&j+3) + b%eval(&j+3) < 4;
%end;

/* No team can have more than four home matches or byes in a row.*/
%do j = 1 %to 14;
    lincon h&j + b&j + h%eval(&j+1) + b%eval(&j+1) + h%eval(&j+2)
           + b%eval(&j+2) + h%eval(&j+3) + b%eval(&j+3) + h%eval(&j+4)
           + b%eval(&j+4) < 5;
%end;
```

```
      /*----------------------------------------------------------*/
      /* Criterion 4 - Weekend Pattern                           */
      /*----------------------------------------------------------*/
      /* Each team plays four weekends at home. */
      lincon 0 %do j = 2 %to 18 %by 2; +h&j %end; =4;
      /* Each team plays four weekends away. */
      lincon 0 %do j = 2 %to 18 %by 2; +a&j %end; =4;
      /* Each team has 1 weekend with a bye  */
      lincon 0 %do j = 2 %to 18 %by 2; +b&j %end; =1;


      /*----------------------------------------------------------*/
      /* Criterion 5 - First Weekends                            */
      /*----------------------------------------------------------*/
      /* Each team must have home games or byes on at least two   */
      /* of the first five weekends.                             */
      lincon 0 %do j = 2 %to 10 %by 2; + h&j + b&j %end; >=2;


      /*----------------------------------------------------------*/
      /* Criterion 9 - (Partial)                                 */
      /*----------------------------------------------------------*/
      /* The team with a bye in date 1 does not play away on the  */
      /*    last date or home in date 17 (Wake)                  */
      /* The team with a bye in date 16 does not play away in    */
      /*    date 18 (Duke)                                       */
      lincon b1 + a18 < 2, b1 + h17 < 2, b16 + a18 < 2;

   run;


   %mend;


   %patterns;
```

The %PATTERNS macro generates 38 patterns. The next step is to find a subset of patterns with cardinality equal to the number of teams that would collectively support a potential assignment to all of the teams. For example, each of the 18 time slots must correspond to four home games, four away games, and one bye. Furthermore, pairs of patterns that do not support a potential meeting date between the two corresponding teams are excluded. The following %PATTERN_SETS macro uses the CLP procedure with the preceding constraints to generate 17 possible pattern sets.

```
/***************************************************************/
/* Determine all possible "pattern sets" considering only time */
/* constraints.                                                */
/* Individual teams are not considered at this stage.          */
/* xi - binary variable indicates pattern i is in pattern set  */
/***************************************************************/
```

```
%macro pattern_sets();

data _null_;
   set all_patterns;
   %do i=1 %to 38;
      if _n_=&i then do;
         %do j=1 %to 18;
            call symput("h&i._&j", put(h&j,best.));
            call symput("a&i._&j", put(a&j,best.));
            call symput("b&i._&j", put(b&j,best.));
         %end;
      end;
   %end;
run;

proc clp out=pattern_sets findall;
   /* xi=1 if pattern i belongs to pattern set */
   var (x1-x38)= [0, 1];

   /* Exactly nine patterns per patterns set */
   lincon 0 %do i = 1 %to 38; + x&i %end;=9;

   /* time slot constraints */
   %do j = 1 %to 18;
      /* Four home games per time slot */
      lincon 0 %do i = 1 %to 38; + &&h&i._&j*x&i %end; =4;
      /* Four away games per time slot */
      lincon 0 %do i = 1 %to 38; + &&a&i._&j*x&i %end; =4;
      /* One bye per time slot */
      lincon 0 %do i = 1 %to 38; + &&b&i._&j*x&i %end; =1;
   %end;

   /* Exclude pattern pairs that do not support a meeting date */
   %do i = 1 %to 38;
      %do i1 = %eval(&i+1) %to 38;
         %let count=0;
         %do j=1 %to 18;
            %if ( (&&h&i._&j=0 or &&a&i1._&j=0) and
                  (&&a&i._&j=0 or &&h&i1._&j=0)) %then %do;
               %let count=%eval(&count+1);
            %end;
         %end;
         %if (&count=18) %then %do;
            lincon x&i+x&i1<=1;
         %end;
      %end;
   %end;
run;

%mend;

%pattern_sets;
```

The %PATTERN_SETS macro generates 17 pattern sets. The final step is to add the individual team constraints and match up teams to the pattern set in order to come up with a schedule for each team. The schedule for each team indicates the opponent for each time slot (0 for a bye) and whether it corresponds to a home game, away game, or a bye.

The following SAS macro %TIMETABLE uses the pattern set index as a parameter and invokes the CLP procedure with the individual team constraints to determine the team schedule.

```
/******************************************************************/
/* Assign individual teams to pattern set k                       */
/* Teams: 1 Clem, 2 Duke, 3 FSU, 4 GT, 5 UMD, 6 UNC, 7 NCSU, 8 UVA, */
/*        9 Wake                                                   */
/******************************************************************/
%macro timetable(k);

proc clp out=ACC_ds_&k varselect=minrmaxc findall;

   %do j = 1 %to 18;
      /*  alpha(i,j): Team i's opponent on date j ( 0 = bye ). */
      %do i = 1 %to 9;
         var alpha&i._&j = [0, 9];
      %end;

      /* Timetable constraint 1 */
      /* Opponents in a time slot must be distinct */
      alldiff ( %do i = 1 %to 9; alpha&i._&j %end; );

      /* Timetable constraint 2 */
      %do i = 1 %to 9;
         %do i1 = 1 %to 9;
            /* indicates if teams i and i1 play in time slot j */
            var X&i._&i1._&j = [0, 1];
            reify X&i._&i1._&j: (alpha&i._&j = &i1);

            /* team i plays i1 iff team i1 plays i */
            %if (&i1 > &i ) %then %do;
               lincon X&i._&i1._&j = X&i1._&i._&j;
            %end;
         %end;
      %end;
   %end;
```

```
/* Mirroring Scheme at team level.                        */
/* The dates are grouped into pairs (j, j1) such that each */
/* team plays the same opponent in dates j and j1.         */
/* One of these should be a home game for each team.        */
%do i = 1 %to 9;
   %do j = 1 %to 18;
      %do j1 = %eval(&j+1) %to 18;
         %if ( &j=1 and &j1=8 ) or ( &j=2 and &j1=9 ) or
             ( &j=3 and &j1=12 ) or ( &j=4 and &j1=13 ) or
             ( &j=5 and &j1=14 ) or ( &j=6 and &j1=15 ) or
             ( &j=7 and &j1=16 ) or ( &j=10 and &j1=17 ) or
             ( &j=11 and &j1=18 ) %then %do;
            lincon alpha&i._&j=alpha&i._&j1,
            /* H and A are matrices that indicate home */
            /* and away games                          */
            H&i._&j=A&i._&j1,
            H&i._&j1=A&i._&j;
         %end;
      %end;
   %end;
%end;


/* Timetable constraint 3 */
/* Each team plays every other team twice */
%do i = 1 %to 9;
   %do i1 = 1 %to 9;
      %if &i1 ne &i %then %do;
         lincon 0 %do j = 1 %to 18; + X&i._&i1._&j %end; = 2;
      %end;
   %end;
%end;


/* Timetable constraint 4 */
/* Teams do not play against themselves */
%do j = 1 %to 18;
   %do i = 1 %to 9;
      lincon alpha&i._&j<>&i;
      lincon X&i._&i._&j = 0;  /* redundant */
   %end;
%end;


/* Timetable constraint 5 */
/* Setup Bye Matrix */
/* alpha&i._&j=0 means team &i has a bye on date &j. */
%do j = 1 %to 18;
   %do i = 1 %to 9;
      var B&i._&j = [0, 1]; /*Bye matrix*/
      reify B&i._&j: ( alpha&i._&j = 0 );
   %end;
%end;
```

```
/* Timetable constraint 6 */
/* alpha&i._&j=&i1 implies teams &i and &i1 play on date &j */
/* It must be a home game for one, away game for the other  */
%do j = 1 %to 18;
   %do i = 1 %to 9;
      %do i1 = 1 %to 9;
         /* reify control variables.*/
         var U&i._&i1._&j = [0, 1] V&i._&i1._&j = [0, 1];

         /* if &i is home and &i1 is away. */
         reify U&i._&i1._&j: ( H&i._&j + A&i1._&j = 2);
         /* if &i1 is home and &i is away. */
         reify V&i._&i1._&j: ( A&i._&j + H&i1._&j = 2);

         /* Necessary condition if &i plays &i1 on date j  */
         lincon X&i._&i1._&j <= U&i._&i1._&j + V&i._&i1._&j;
      %end;
   %end;
%end;


/* Timetable constraint 7 */
/* Each team must be home, away or have a bye on a given date */
%do j = 1 %to 18;
   %do i = 1 %to 9;
      /* Team &i is home (away) at date &j. */
      var H&i._&j = [0, 1] A&i._&j = [0, 1];
      lincon H&i._&j + A&i._&j + B&i._&j = 1;
   %end;
%end;

%do i = 1 %to 9;
   %do i1 = %eval(&i+1) %to 9;

      /* Timetable constraint 8 */
      /*-----------------------------------------------------*/
      /* Criterion 6 - Rival Matches                         */
      /*-----------------------------------------------------*/
      /* The final weekend is reserved for 'rival games' */
      /* unless the team plays FSU or has a bye          */
      %if ( &i=1 and &i1=4 ) or ( &i=2 and &i1=6 ) or
          ( &i=5 and &i1=8 ) or ( &i=7 and &i1=9 ) %then %do;
         lincon X&i._&i1._18 + B&i._18 + X&i._3_18 = 1;

         /* redundant */
         lincon X&i1._&i._18 + B&i1._18 + X&i1._3_18 = 1;
      %end;
```

```
          /* Timetable constraint 9 */
          /*---------------------------------------------------------*/
          /* Criterion 7 - Popular Matches                           */
          /*---------------------------------------------------------*/
          /* The following pairings are specified to occur at        */
          /* least once in February.                                 */
          %if ( &i=2 and &i1=4 ) or ( &i=2 and &i1=9 ) or
              ( &i=4 and &i1=6 ) or ( &i=6 and &i1=9 ) %then %do;
             lincon 0 %do j = 11 %to 18; + X&i._&i1._&j  %end; >= 1;

             /* redundant */
             lincon 0 %do j = 11 %to 18; + X&i1._&i._&j %end; >= 1;
          %end;
       %end;
   %end;

   /* Timetable constraint 10 */
   /*---------------------------------------------------------*/
   /* Criterion 8 - Opponent Sequence                         */
   /*---------------------------------------------------------*/
   %do i = 1 %to 9;
      /* No team plays two consecutive away dates against    */
      /* Duke (2) and UNC (6)                                */
      %do j = 1 %to 17;
         var Q&i._26_&j = [0, 1] P&i._26_&j = [0, 1];
         reify Q&i._26_&j: ( X&i._2_&j +  X&i._6_&j = 1 );
         reify P&i._26_&j: ( X&i._2_%eval(&j+1) +  X&i._6_%eval(&j+1) =1 );
         lincon Q&i._26_&j + A&i._&j + P&i._26_&j + A&i._%eval(&j+1) < 4;
      %end;

      /* No team plays three consecutive dates against       */
      /* Duke(2), UNC(6) and Wake(9).                        */
      %do j = 1 %to 16;
         var L&i._269_&j = [0, 1] M&i._269_&j = [0, 1]
             N&i._269_&j = [0, 1];
         reify L&i._269_&j: ( X&i._2_&j +  X&i._6_&j +  X&i._9_&j = 1 );
         reify M&i._269_&j: ( X&i._2_%eval(&j+1) + X&i._6_%eval(&j+1) +
                              X&i._9_%eval(&j+1) =1 );
         reify N&i._269_&j: ( X&i._2_%eval(&j+2) + X&i._6_%eval(&j+2) +
                              X&i._9_%eval(&j+2) =1 );
         lincon L&i._269_&j + M&i._269_&j + N&i._269_&j  < 3;
      %end;
   %end;
```

```
   /* Timetable constraint 11 */
   /*----------------------------------------------------*/
   /* Criterion 9 - Idiosyncratic Constraints           */
   /*----------------------------------------------------*/
   /* UNC plays Duke in date 11 and 18 */
   lincon alpha6_11 = 2 ;
   lincon alpha6_18 = 2 ;
   /* UNC plays Clem in the second date. */
   lincon alpha6_2 = 1 ;
   /* Duke has a bye in date 16. */
   lincon B2_16 = 1 ;
   /* Wake does not play home in date 17. */
   lincon H9_17 = 0 ;
   /* Wake has a bye in the first date. */
   lincon B9_1 = 1 ;
   /* Clem, Duke, UMD and Wake do not play away in the last date. */
   lincon A1_18 = 0 ;
   lincon A2_18 = 0 ;
   lincon A5_18 = 0 ;
   lincon A9_18 = 0 ;
   /* Clem, FSU, and GT do not play away in the first date. */
   lincon A1_1 = 0 ;
   lincon A3_1 = 0 ;
   lincon A4_1 = 0 ;
   /* FSU and NCSU do not have a bye in the last date. */
   lincon B3_18 = 0 ;
   lincon B7_18 = 0 ;
   /* UNC does not have a bye in the first date. */
   lincon B6_1 = 0 ;

   /* Timetable constraint 12 */
   /*----------------------------------------------------*/
   /* Match teams with patterns.                         */
   /*----------------------------------------------------*/
   %do i = 1 %to 9; /* For each team */
      var p&i=[1,9];
      %do j=1 %to 18; /* For each date */
         element ( p&i, (&&col&k._h_&j), H&i._&j )
                 ( p&i, (&&col&k._a_&j), A&i._&j )
                 ( p&i, (&&col&k._b_&j), B&i._&j );
      %end;
   %end;
run;

%mend;
```

```
/**********************************************************/
/* Try all possible pattern sets to find all valid schedules. */
/**********************************************************/

%macro find_schedules;

proc transpose data=pattern_sets out=trans_good; run;

data _temp;
   set trans_good;
   set all_patterns;
run;

proc sql noprint;
   %do k = 1 %to 17; /* For each pattern */
      %do j=1 %to 18; /* For each date */
         select h&j into :col&k._h_&j
            separated by ',' from _temp where col&k=1;
         select a&j into :col&k._a_&j
            separated by ',' from _temp where col&k=1;
         select b&j into :col&k._b_&j
            separated by ',' from _temp where col&k=1;
      %end;
   %end;
run;

data all; run;

%do k = 1 %to 17; /* For each pattern set */
   %timetable(k=&k);

   data all;
      set all ACC_ds_&k;
   run;
%end;

data all;
   set all;
   if _n_=1 then delete;
run;

%mend;

%find_schedules;
```

The %FIND_SCHEDULES macro invokes the %TIMETABLE macro for each of the 17 pattern sets and generates 179 possible schedules including the one that was eventually used by the ACC, which is displayed in Output 3.12.1.

Output 3.12.1   ACC Basketball Tournament Schedule

ACC Basketball Tournament Scheduling

Team	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
Clem	UMD	UNC	@Wake		FSU	@Duke	UVA	@UMD	@UNC	NCSU	@GT	Wake		@FSU	Duke	@UVA	@NCSU	GT
Duke	UVA	@UMD	NCSU	@FSU	@Wake	Clem		@UVA	UMD	GT	@UNC	@NCSU	FSU	Wake	@Clem		@GT	UNC
FSU	UNC	@NCSU	@UMD	Duke	@Clem	@GT	Wake	@UNC	NCSU		@UVA	UMD	@Duke	Clem	GT	@Wake	UVA	
GT	NCSU		@UNC	Wake	@UVA	FSU	UMD	@NCSU		@Duke	Clem	UNC	@Wake	UVA	@FSU	@UMD	Duke	@Clem
UMD	@Clem	Duke	FSU	@NCSU	UNC	@Wake	@GT	Clem	@Duke	UVA		@FSU	NCSU	@UNC	Wake	GT		@UVA
UNC	@FSU	@Clem	GT	UVA	@UMD		@NCSU	FSU	Clem	@Wake	Duke	@GT	@UVA	UMD	NCSU	@Duke	Wake	
NCSU	@GT	FSU	@Duke	UMD		UVA	UNC	GT	@FSU	@Clem	@Wake	Duke	@UMD		@UNC	Clem	Wake	@UVA
UVA	@Duke	Wake		@UNC	GT	@NCSU	@Clem	Duke	@Wake	@UMD	FSU		UNC	@GT	NCSU	Clem	@FSU	UMD
Wake		@UVA	Clem	@GT	Duke	UMD	@FSU		UVA	UNC	NCSU	@Clem	GT	@Duke	@UMD	FSU	@UNC	NCSU

References

Applegate, D. and Cook, W. (1991), "A Computational Study of the Job Shop Scheduling Problem," *ORSA Journal on Computing*, 3, 149–156.

Baptiste, P. and Le Pape, C. (1996), "Edge-Finding Constraint Propagation Algorithms for Disjunctive and Cumulative Scheduling," in *Proceedings of the 15th Workshop of the UK Planning Special Interest Group*, Liverpool, UK.

Bartusch, M. (1983), *Optimierung von Netzplänen mit Anordnungsbeziehungen bei knappen Betriebsmitteln*, Ph.D. thesis, Universität Passau, Fakultät für Mathematik und Informatik.

Carlier, J. and Pinson, E. (1989), "An Algorithm for Solving the Job-Shop Scheduling Problem," *Management Science*, 35(2), 164–176.

Carlier, J. and Pinson, E. (1990), "A Practical Use of Jackson's Preemptive Schedule for Solving the Job-Shop Problem," *Annals of Operations Research*, 26, 269–287.

Colmerauer, A. (1990), "An Introduction to PROLOG III," *Communications of the ACM*, 33(7), 70–90.

Dincbas, M., Simonis, H., and Van Hentenryck, P. (1988), "Solving the Car-Sequencing Problem in Constraint Logic Programming," in Y. Kodratoff, ed., *Proceedings of ECAI-88*, 290–295, Munich, W. Germany.

Floyd, R. W. (1967), "Nondeterministic Algorithms," *Journal of the ACM*, 14(4), 636–644.

Garey, M. R. and Johnson, D. S. (1979), *Computers and Intractability: A Guide to the Theory of NP-Completeness*, New York: W. H. Freeman & Co.

Gravel, M., Gagne, C., and Price, W. L. (2005), "Review and Comparison of Three Methods for the Solution of the Car Sequencing Problem," *Journal of the Operational Research Society*, 56, 1287–1295.

Haralick, R. M. and Elliot, G. L. (1980), "Increasing Tree Search Efficiency for Constraint Satisfaction Problems," *Artificial Intelligence*, 14(3), 263–313.

Henz, M. (2001), "Scheduling a Major College Basketball Conference—Revisited," *Operations Research*, 49, 163–168.

Jaffar, J. and Lassez, J. (1987), "Constraint Logic Programming," in *Proceedings of the 14th Annual ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*, 111–119, Munich, W. Germany.

Kumar, V. (1992), "Algorithms for Constraint-Satisfaction Problems: A Survey," *AI Magazine*, 13, 32–44.

Lawrence, S. (1984), *Resource Constrained Project Scheduling: An Experimental Investigation of Heuristic Scheduling Techniques (Supplement)*, Technical report, Graduate School of Industrial Administration, Carnegie Mellon University, Pittsburgh, PA.

Mackworth, A. K. (1977), "Consistency in Networks of Relations," *Artificial Intelligence*, 8, 99–118.

Muth, J. F. and Thompson, G. L., eds. (1963), *Industrial Scheduling*, Englewood Cliffs, NJ: Prentice Hall.

Nemhauser, G. L. and Trick, M. A. (1998), "Scheduling a Major College Basketball Conference," *Operations Research*, 46, 1–8.

Nemhauser, G. L. and Wolsey, L. A. (1988), *Integer and Combinatorial Optimization*, New York: John Wiley & Sons.

Nuijten, W. (1994), *Time and Resource Constrained Scheduling*, Ph.D. thesis, Eindhoven Institute of Technology, Eindhoven, Netherlands.

Riley, P. and Taalman, L. (2008), "Brainfreeze Puzzles," .
  URL http://www.brainfreezepuzzles.com/main/piday2008.html

Smith, B. M., Brailsford, S. C., Hubbard, P. M., and Williams, H. P. (1996), "The Progressive Party Problem: Integer Linear Programming and Constraint Programming Compared," *Constraints*, 1, 119–138.

Sokol, J. (2002), *Modeling Automobile Paint Blocking: A Time Window Traveling Salesman Problem*, Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA.

Solnon, C., Cung, V. D., Nguyen, A., and Artigues, C. (2008), "The Car Sequencing Problem: Overview of State-of-the-Art Methods and Industrial Case-Study of the ROADEF 2005 Challenge Problem," *European Journal of Operational Research*, 191(3), 912–927.

Trick, M. (2004), "Constraint Programming: A Tutorial," http://mat.gsia.cmu.edu/trick/cp.ppt.

Tsang, E. (1993), *Foundations of Constraint Satisfaction*, London: Academic Press.

Van Hentenryck, P. (1989), *Constraint Satisfaction in Logic Programming*, Cambridge, MA: MIT Press.

Van Hentenryck, P. (2002), "Constraint and Integer Programming in OPL," *INFORMS Journal on Computing*, 14(4), 345–372.

Van Hentenryck, P., Deville, Y., and Teng, C. (1992), "A Generic Arc-Consistency Algorithm and Its Specializations," *Artificial Intelligence*, 57(2–3), 291–321.

Waltz, D. L. (1975), "Understanding Line Drawings of Scenes with Shadows," in P. H. Winston, ed., *The Psychology of Computer Vision*, 19–91, New York: McGraw-Hill.

Williams, H. P. and Wilson, J. M. (1998), "Connections between Integer Linear Programming and Constraint Logic Programming—An Overview and Introduction to the Cluster of Articles," *INFORMS Journal of Computing*, 10(3), 261–264.

Subject Index

Syntax Index

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**.