

# Chapter 1

## Introduction to Optimization

### Chapter Table of Contents

---

<b>OVERVIEW</b> . . . . .	3
<b>DATA FLOW</b> . . . . .	4
PROC LP . . . . .	4
PROC NETFLOW . . . . .	5
PROC NLP . . . . .	8
PROC TRANS . . . . .	11
PROC ASSIGN . . . . .	12
Model Formats: PROC LP and PROC NETFLOW . . . . .	12
PROC ASSIGN and PROC TRANS Data Formats . . . . .	23
<b>MODEL BUILDING</b> . . . . .	24
<b>MATRIX GENERATION</b> . . . . .	26
Exploiting Model Structure . . . . .	28
<b>REPORT WRITING</b> . . . . .	31
The DATA step . . . . .	31
Other Reporting Procedures . . . . .	32
<b>DECISION SUPPORT SYSTEMS</b> . . . . .	34
The Full-Screen Interface . . . . .	34
Communicating with the Optimization Procedures . . . . .	35
<b>REFERENCES</b> . . . . .	35



# Chapter 1

## Introduction to Optimization

---

### Overview

This chapter describes how to use SAS/OR software to solve a wide variety of optimization problems. The basic optimization problem is that of minimizing or maximizing an objective function subject to constraints imposed on the variables of that function. The objective function and constraints can be linear or nonlinear; the constraints can be bound constraints, equality or inequality constraints, or integer constraints.

Traditionally, optimization problems are divided into Linear Programming (LP; all functions are linear) and Nonlinear Programming (NLP). Variations of LP problems are assignment problems, network flow problems, and transportation problems. Nonlinear Regression (fitting a nonlinear model to a set of data and the subsequent statistical analysis of the results) is a special NLP problem. Since these applications are so common, SAS/OR software has separate procedures or facilities within procedures for solving each type of these problems. Model data are supplied in a form suited for the particular type of problem. Another benefit is that an optimization algorithm can be specialized for the particular type of problem, reducing solution times. Optimizers can exploit some structure in problems such as imbedded networks, special ordered sets, least squares, and quadratic objective functions.

SAS/OR software has five procedures used for optimization:

- **PROC ASSIGN** for solving assignment problems
- **PROC LP** for solving linear and mixed integer programming problems
- **PROC NETFLOW** for network programming problems with side constraints, and linear programming problems solved by an Interior Point algorithm
- **PROC NLP** for non-linear programming problems
- **PROC TRANS** for solving transportation problems

SAS/OR procedures use syntax that is similar to other SAS procedures. In particular, all SAS retrieval, data management, reporting and analysis can be used with SAS/OR software. Each optimizer is designed to integrate with the SAS System to simplify model building, maintenance, solution, and report writing.

Data for models are supplied to SAS/OR procedures in SAS data sets. These data sets can be saved and easily changed and the problem can be re-solved. Because the models are in SAS data sets, problem data that can represent pieces of a larger model can be concatenated and merged. The SAS/OR procedures output SAS data sets containing the solutions. These can then be used to produce customized reports.

This structure allows decision support systems to be constructed using SAS/OR procedures and other tools in the SAS System as building blocks.

The following list suggests application areas where decision support systems have been used. In practice, models often contain elements of several applications listed here.

- **Product-Mix problems** find the mix of products that generates the largest return when there are several products that compete for limited resources.
- **Blending problems** find the mix of ingredients to be used in a product so that it meets minimum standards at minimum cost.
- **Time-Staged problems** are models whose structure repeats as a function of time. Production and inventory models are classic examples of time-staged problems. In each period, production plus inventory minus current demand equals inventory carried to the next period.
- **Scheduling problems** assigns people to times, places, or tasks so as to optimize peoples' preferences while satisfying the demands of the schedule.
- **Multiple objective problems** have multiple conflicting objectives. Typically, the objectives are prioritized and the problems are solved sequentially in a priority order.
- **Capital budgeting and project selection problems** ask for the project or set of projects that will yield the greatest return.
- **Location problems** seek the set of locations that meets the distribution needs at minimum cost.
- **Cutting stock problems** find the partition of raw material that minimizes waste.

---

## Data Flow

The procedures ASSIGN, LP, NETFLOW, NLP, and TRANS take a model that has been saved in one or more SAS data sets, solve it, and save the solution in other SAS data sets. Each procedure defines a SAS macro variable that contains a character string indicating whether or not the procedure terminated successfully and the status of the optimizer (for example, whether the optimum was found). This information is useful when the procedure is one of the steps in a larger program.

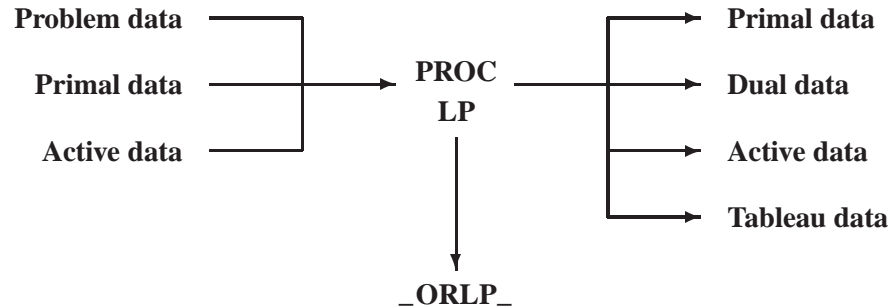
---

## PROC LP

The LP procedure solves linear and mixed integer programs. It can perform several types of post-optimality analysis, including range analysis, sensitivity analysis, and parametric programming. The procedure can also be used interactively.

PROC LP requires a problem data set that contains the model. In addition, a primal and active data set can be used for warm starting a problem that has been partially solved previously.

The following diagram illustrates all the input and output data sets that are possible with PROC LP. It also shows the macro variable `_ORLP_` that PROC LP defines.



**Figure 1.1.** Input and Output Data Sets in PROC LP

The problem data describing the model can be in one of two formats: a sparse or a dense format. The dense format represents the model as a rectangular matrix. The sparse format represents only the nonzero elements of a rectangular matrix. The sparse and dense input formats are described in more detail later in this chapter.

---

## PROC NETFLOW

Constrained network models describe a wide variety of real-world applications ranging from production, inventory, and distribution problems to financial applications. These problems can be solved with the NETFLOW procedure.

These models are conceptionally easy since they are based on network diagrams that represent the problem pictorially. This procedure accepts the network specification in a format that is particularly suited to networks. This not only simplifies problem description but also aids in the interpretation of the solution.

Certain algebraic features of networks are exploited by a specialized version of the Simplex method so that solution times are reduced. Another class of optimization algorithm, the Interior Point algorithm, has been implemented in PROC NETFLOW and can be used as an alternative to the Simplex algorithm to solve network problems.

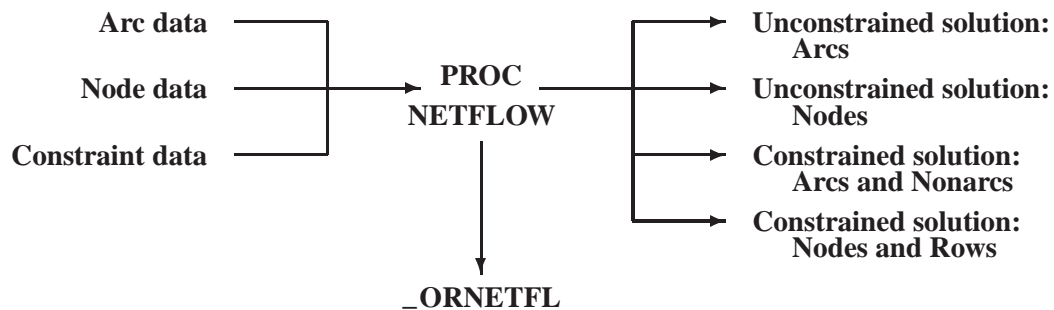
Should NETFLOW detect there are no arcs and nodes in the model's data, (that is, there is no network component), it solves a linear programming (LP) problem. The Interior Point algorithm is automatically elected to perform the optimization. NETFLOW's ability to solve LP problems is an important ability new for the Version 7 release of the SAS System.

### **Networks and the Simplex Network Algorithm**

PROC NETFLOW's network Simplex algorithm solves pure network flow problems and network flow problems with linear side constraints. The side constraints can have variables directly related to the network (called *arc variables*) as well as variables (called *nonarc variables*) that have nothing to do with the network. The procedure accepts the network specification in a format that is particularly suited to networks.

Although network problems could be solved by PROC LP, the NETFLOW procedure generally solves network flow problems more efficiently than PROC LP.

Network flow problems, such as finding the minimum cost flow in a network, require model representation in a format that is simpler than PROC LP. The network is represented in two data sets: a node data set that names the nodes in the network and gives supply and demand information at them, and an arc data set that defines the arcs in the network using the node names and gives arc costs and capacities. In addition, a side-constraint data set is included that gives any side constraints that apply to the flow through the network. Examples of these are found later in this chapter.



**Figure 1.2.** Input and Output Data Sets: PROC NETFLOW: Simplex Algorithm

The constraint data contains constraints that are not network flow conservation constraints. These can be specified in either the sparse or dense input formats. This is the same format that is used by PROC LP so that any of the model building techniques that apply to models for PROC LP also apply for network flow models having side constraints.

The NETFLOW procedure saves solutions in four data sets. Two of these store solutions for the pure network model, ignoring the restrictions imposed by the side constraints. The remaining two data sets contain the solutions to the network flow problem when the side constraints apply.

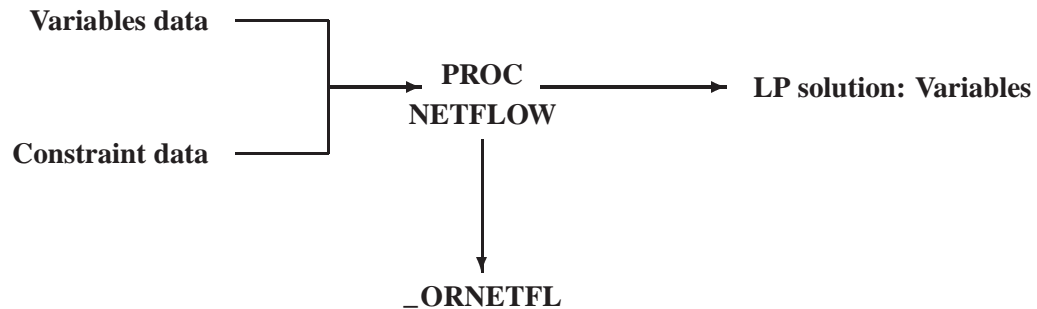
### **Network and Linear Programs solved by the Interior Point algorithm**

The Simplex algorithm, developed shortly after World War II, was the main method used to solve Linear Programming problems. Over the last decade, the Interior Point algorithm has been developed to also solve Linear Programming problems. From the start it showed great theoretical promise, and considerable research in the area resulted in practical implementations that performed competitively with the Simplex algorithm. More recently, Interior Point algorithms have evolved to become superior to the Simplex algorithm, in general, especially when the problems are large.

The Interior Point algorithm has been implemented in PROC NETFLOW. This algorithm solves Linear Programs, as well as network problems. When PROC NETFLOW detects that the problem has no network component, it automatically invokes the Interior Point algorithm to solve the problem.

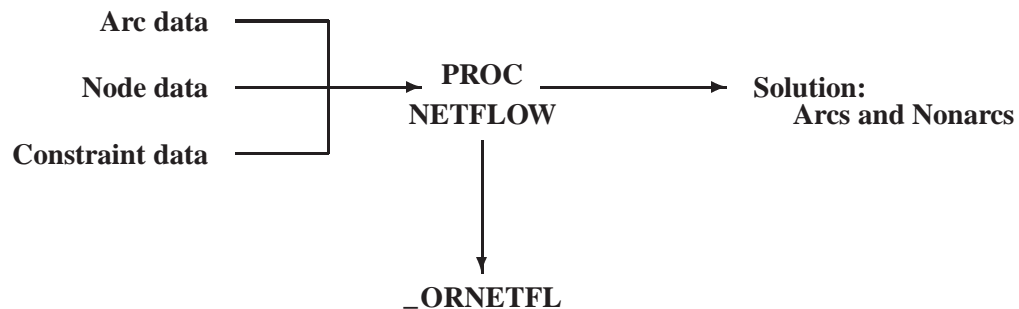
The data required by PROC NETFLOW for a Linear Program resembles the data for nonarc variables and constraints for constrained network problems. It also is similar to the data required by PROC LP.

The LP is represented in one or two data sets: a data set that defines the variables in the LP using variable names and gives objective function coefficients and upper and lower bounds. In addition, a constraint data set can be included that gives any constraints.



**Figure 1.3.** Input and Output Data Sets: PROC NETFLOW: LP problems

When solving a constrained network problem, indicate that the Interior Point algorithm is to be used by specifying the INTPOINT option. The input data is the same whether the Simplex or Interior Point method is used. In other aspects of how PROC NETFLOW is used, there are only minor changes. The output, particularly the CONOUT data set where the solution is sent, is similar no matter which optimizer is used. The Interior Point method is often faster when problems have many side constraints.



**Figure 1.4.** Input and Output Data Sets: PROC NETFLOW: Network problems: Intpoint algorithm

The constraint data can be specified in either the sparse or dense input format. This is the same format that is used by PROC LP so that any of the model building techniques that apply to models for PROC LP also apply for LP models solved by PROC NETFLOW.

---

## PROC NLP

The NLP procedure (**NonLinear Programming**) offers a set of optimization techniques for minimizing or maximizing a continuous nonlinear function  $f(x)$  of  $n$  decision variables,  $x = (x_1, \dots, x_n)^T$  with lower and upper bound, linear and nonlinear, equality and inequality constraints. This can be expressed as solving

$$\begin{array}{ll} \min_{x \in \mathcal{R}^n} & f(x) \\ \text{subject to} & c_i(x) = 0 \quad i = 1, \dots, m_e \\ & c_i(x) \geq 0 \quad i = m_e + 1, \dots, m \\ & u_i \geq x_i \geq l_i \quad i = 1, \dots, n \end{array}$$

where  $f$  is the objective function, the  $c_i$ 's are the constraint functions, and  $u_i, l_i$ 's are the upper and lower bounds. Problems of this type are found in many settings ranging from optimal control to maximum likelihood estimation.

The NLP procedure provides a number of algorithms for solving this problem that take advantage of special structure on the objective function or constraints, or both. One example is the **quadratic programming problem**:

$$\begin{array}{ll} f(x) = \frac{1}{2}x^T Gx + g^T x + b \\ \text{subject to} & c_i(x) = 0 \quad i = 1, \dots, m_e \end{array}$$

where the  $c_i(x)$ 's are linear functions;  $g = (g_1, \dots, g_n)^T$  and  $b = (b_1, \dots, b_n)^T$  are vectors; and  $G$  is an  $n \times n$  symmetric matrix.

Another example is the **least-squares problem**:

$$\begin{array}{ll} f(x) = \frac{1}{2}\{f_1^2(x) + \dots + f_l^2(x)\} \\ \text{subject to} & c_i(x) = 0 \quad i = 1, \dots, m_e \end{array}$$

where the  $c_i(x)$ 's are linear functions, and  $f_1(x), \dots, f_m(x)$  are nonlinear functions of  $x$ .

The following optimization techniques are supported in PROC NLP:

- Quadratic Active Set Technique
- Trust-Region Method
- Newton-Raphson Method With Line Search
- Newton-Raphson Method With Ridging
- Quasi-Newton Methods
- Double-Dogleg Method
- Conjugate Gradient Methods
- Nelder-Mead Simplex Method



- Levenberg-Marquardt Method
- Hybrid Quasi-Newton Methods

These optimization techniques require a continuous objective function  $f$ , and all but one (NMSIMP) require continuous first-order derivatives of the objective function  $f$ . Some of the techniques also require continuous second-order derivatives. There are three ways to compute derivatives in PROC NLP:

1. analytically (using a special derivative compiler), the default method
2. via finite difference approximations
3. via user-supplied exact or approximate numerical functions


Nonlinear programs can be input into the procedure in various ways. The objective, constraint, and derivative functions are specified using the programming statements of PROC NLP. In addition, information in SAS data sets can be used to define the structure of objectives and constraints as well as specify constants used in objectives, constraints, and derivatives.

PROC NLP uses data sets to input various pieces of information.

- The DATA= data set enables you to specify data shared by all functions involved in a least squares problem.
- The INQUAD= data set contains the arrays appearing in a quadratic programming problem.
- The INVAR= data set specifies initial values for the decision variables, the values of constants that are referred to in the program statements, and simple boundary and general linear constraints.
- The MODEL= data set specifies a model (functions, constraints, derivatives) saved at a previous execution of the NLP procedure.

PROC NLP uses data sets to output various results.

- The OUTVAR= data set saves the values of the decision variables, the derivatives, the solution, and the covariance matrix at the solution.
- The OUT= output data set contains variables generated in the program statements defining the objective function as well as selected variables of the DATA= input data set, if available.
- The OUTMODEL= data set saves the programming statements. It can be used to input a model in the MODEL= input data set.



**Figure 1.5.** Input and Output Data Sets in PROC NLP

As an alternative to supplying data in SAS data sets, some or all data for the model can be specified using SAS programming statements. These are similar to those used in the SAS DATA step.

Consider the simple example of minimizing the Rosenbrock Function (Rosenbrock 1960).

$$\begin{aligned} f(x) &= \frac{1}{2}\{100(x_2 - x_1^2)^2 + (1 - x_1)^2\} \\ &= \frac{1}{2}\{f_1^2(x) + f_2^2(x)\}, \quad x = (x_1, x_2) \end{aligned}$$

The minimum function value is  $f(x^*) = 0$  at  $x^* = (1, 1)$ . This problem does not have any constraints.

The following PROC NLP run be used to solve this problem:

```
proc nlp;
  min f;
  decvar x1 x2;
  f1 = 10 * (x2 - x1 * x1);
  f2 = 1 - x1;
  f = .5 * (f1 * f1 + f2 * f2);
run;
```

The MIN statement identifies the symbol  $f$  that characterizes the objective function in terms of  $f_1$  and  $f_2$ , and the DECVAR statement names the decision variables  $X_1$  and  $X_2$ . Because there is no explicit optimizing algorithm option specified (TECH=), PROC NLP would use the Newton-Raphson method with ridging, the default algorithm when there are no constraints.

A better way to solve this problem is to take advantage of the fact that  $f$  is a sum of squares of  $f_1$  and  $f_2$  and to treat it as a least-squares problem. Using the LSQ statement instead of the MIN statement tells the procedure that this is a least-squares

problem, which results in the use of one of the specialized algorithms for solving least-squares problems (for example, Levenberg-Marquardt).

```
proc nlp;
  lsq f1 f2;
  decvar x1 x2;
  f1 = 10 * (x2 - x1 * x1);
  f2 = 1 - x1;
run;
```

The LSQ statement results in the minimization of a function that is the sum of squares of functions that appear in the LSQ statement.

The least-squares specification is preferred because it enables the procedure to exploit the structure in the problem for numeric stability and performance.

There are several other NLP statements that are used to supply additional data of the model, such as variable value bounds and linear and nonlinear constraints. The following is an example of a problem with bounds and with linear and nonlinear constraints:

```
proc nlp tech=QUANEW;
  min f;
  decvar x1 x2;
  bounds x1 - x2 <= .5;
  lincon x1 + x2 <= .6;
  nlincon c1 >= 0;

  c1 = x1 * x1 - 2 * x2;


  f1 = 10 * (x2 - x1 * x1);
  f2 = 1 - x1;

  f = .5 * (f1 * f1 + f2 * f2);
run;
```

---

## PROC TRANS

Transportation networks are a special type of network, called *bipartite* networks, that have only supply and demand nodes and arcs directed from supply nodes to demand nodes. For these networks, data can be given most efficiently in a rectangular or matrix form. The TRANS procedure takes cost, capacity, and lower bound data in this form. The observations in these data sets correspond to supply nodes, and the variables correspond to demand nodes.




**Figure 1.6.** Input and Output Data Sets in PROC TRANS

The TRANS procedure puts the solution in a single output data set. The details are described in the chapter on the TRANS procedure. As with the other optimization procedures, the TRANS procedure defines a macro variable, named `_ORTRANS`, that has a character string that describes the solution status.

---

## PROC ASSIGN

The assignment problem is a special type of transportation problem, one having supply and demand values of one unit. As with the transportation problem, the cost data for this type of problem are saved in a SAS data set in rectangular form.



**Figure 1.7.** Input and Output Data Sets in PROC ASSIGN

The ASSIGN procedure also saves the solution in a SAS data set and defines a macro variable.

---

## Model Formats: PROC LP and PROC NETFLOW

Model generation and maintenance are often difficult and expensive aspects of applying mathematical programming techniques. The flexible input formats for the optimization procedures in SAS/OR software simplify this task.

A small product mix problem serves as a starting point for a discussion of different types of model formats supported in SAS/OR software.

A candy manufacturer makes two products: chocolates and toffee. What combination of chocolates and toffee should be produced in a day in order to maximize the company's profit? Chocolates contribute \$0.25 per pound to profit and toffee contributes \$0.75 per pound. The variables chocolates and toffee are the decision variables.

Four processes are used to manufacture the candy:

1. Process 1 combines and cooks the basic ingredients for both chocolates and toffee.
2. Process 2 adds colors and flavors to the toffee, then cools and shapes the confection.
3. Process 3 chops and mixes nuts and raisins, adds them to the chocolates, then cools and cuts the bars.
4. Process 4 is packaging: chocolates are placed in individual paper shells; toffee are wrapped in cellophane packages.

During the day, there are 7.5 hours (27,000 seconds) available for each process.

Firm time standards have been established for each process. For Process 1, mixing and cooking take 15 seconds for each pound of chocolate and 40 seconds for each pound of toffee. Process 2 takes 56.25 seconds per pound of toffee. For Process 3, 18.75 seconds are required for each pound of chocolate. In packaging, a pound of chocolates can be wrapped in 12 seconds, whereas 50 seconds are required for a pound of toffee. These data are summarized below:

Process	Available	Required per Pound	
	Time (sec)	chocolates (sec)	toffee (sec)
1 Cooking	27,000	15	40
2 Color/Flavor	27,000		56.25
3 Condiments	27,000	18.75	
4 Packaging	27,000	126	50

The object is to

$$\text{Maximize: } 0.25(\text{chocolates}) + 0.75(\text{toffee})$$

which is the company's total profit.

The production of the candy is limited by the time available for each process. The limits placed on production by Process 1 are expressed by the inequality

$$\text{Process 1: } 15(\text{chocolates}) + 40(\text{toffee}) \leq 27,000$$

Process 1 can handle any combination of chocolates and toffee that satisfies this inequality.

The limits on production by other processes generate constraints described by the following inequalities.

Process 2:  $56.25(\text{toffee}) \leq 27,000$   
 Process 3:  $18.75(\text{chocolates}) \leq 27,000$   
 Process 4:  $12(\text{chocolates}) + 50(\text{toffee}) \leq 27,000$

This linear program illustrates the type of problem known as a product mix example. The mix of products that maximizes the objective without violating the constraints is the solution. There are two formats that can represent this model.

### Dense format

The following DATA step creates a SAS data set for the preceding problem. Notice that the values of CHOCO and TOFFEE in the data set are the coefficients of those variables in the equations corresponding to the objective function and constraints. The variable `_id_` contains a character string that names the rows in the problem data set. The variable `_type_` is a character variable that contains keywords that describes the type of each row in the problem data set. And the variable `_rhs_` contains the right-hand-side values.

```
data factory;
  input _id_ $ CHOCO TOFFEE _type_ $ _rhs_;
  datalines;
object      0.25    0.75    MAX      .
process1    15.00   40.00   LE      27000
process2     0.00   56.25   LE      27000
process3    18.75    0.00   LE      27000
process4    12.00   50.00   LE      27000
  ;
```

To solve this using the Interior Point algorithm of PROC NETFLOW, specify:

```
proc netflow arcdata=factory condata=factory;
```

However, this example will be solved by the LP procedure. Because the special variables `_id_`, `_type_`, and `_rhs_` are used in the problem data set, there is no need to identify them to the LP procedure.

```
proc lp;
```

The output from the LP procedure is in four sections.

### Problem Summary

The first section of the output, PROBLEM SUMMARY, describes the problem by identifying the objective function (defined by the first observation in the data set used as input); the right-hand-side variable; the type variable; and the density of the problem. The problem density describes the relative number of elements in the problem matrix that are nonzero. The fewer zeros in the matrix, the higher the problem density. The PROBLEM SUMMARY describes the problem giving the number and type of variables in the model and the number and type of constraints. The types of variables in the problem are identified. Variables are either structural or logical. Structural

variables are identified in the VAR statement, when the dense format is used. They are the unknowns in the equations defining the objective function and constraints. By default, PROC LP assumes that structural variables have the additional constraint that they must be nonnegative. Upper and lower bounds to structural variables can be defined.

The SAS System	
The LP Procedure	
Problem Summary	
Objective Function	Max object
Rhs Variable	_rhs_
Type Variable	_type_
Problem Density (%)	41.67
Variables	Number
Non-negative	2
Slack	4
Total	6
Constraints	Number
LE	4
Objective	1
Total	5

**Figure 1.8.** PROBLEM SUMMARY

The Problem Summary shows, for example, that there are 2 nonnegative decision variables, namely CHOCO and TOFFEE. It also shows that there are 4 constraints of type LE.

After the procedure displays this information, it solves the problem, then produces the SOLUTION SUMMARY.

### **Solution Summary**

This section (Figure 1.9) gives information about the solution that was found; whether the optimizer terminated successfully having found the optimum.

When PROC LP solves a problem, an iterative process is used. First, the procedure finds a feasible solution, one that satisfies the constraints. The second phase finds the optimum solution from the set of feasible solutions. The SOLUTION SUMMARY gives the number of iterations the procedure took in each of these phases, the number of variables in the initial feasible solution, the time the procedure used to solve the problem, and the number of matrix inversions necessary.

The LP Procedure	
Solution Summary	
Terminated Successfully	
Objective Value	475
Phase 1 Iterations	0
Phase 2 Iterations	3
Phase 3 Iterations	0
Integer Iterations	0
Integer Solutions	0
Initial Basic Feasible Variables	6
Time Used (seconds)	0
Number of Inversions	2
Epsilon	1E-8
Infinity	1.797693E308
Maximum Phase 1 Iterations	100
Maximum Phase 2 Iterations	100
Maximum Phase 3 Iterations	99999999
Maximum Integer Iterations	100
Time Limit (seconds)	120

**Figure 1.9.** SOLUTION SUMMARY

After performing 3 phase 2 iterations, the procedure terminated successfully finding an optimum with optimal objective value of 475.

### Variable Summary

The next section of the output is the VARIABLE SUMMARY. For each variable, the variable summary (Figure 1.10) gives the value, objective function coefficient, status in the solution, and reduced cost.

The LP Procedure						
Variable Summary						
Variable	Col Name	Status	Type	Price	Activity	Reduced Cost
1 CHOCO		BASIC	NON-NEG	0.25	1000	0
2 TOFFEE		BASIC	NON-NEG	0.75	300	0
3 process1			SLACK	0	0	-0.012963
4 process2		BASIC	SLACK	0	10125	0
5 process3		BASIC	SLACK	0	8250	0
6 process4			SLACK	0	0	-0.00463

**Figure 1.10.** VARIABLE SUMMARY

The VARIABLE SUMMARY contains details about each variable in the solution. The Activity column shows that optimum profitability is achieved when 1000 pounds of chocolate and 300 pounds of toffee are produced. The variables process1, process2, process3, and process4 correspond to the four slack variables in the Process 1, Process 2, Process 3, and Process 4 constraints, respectively. Producing 1000 pounds of chocolate and 300 pounds of toffee a day leaves 10,125 seconds of slack time in Process 2 (where colors and flavors are added to the toffee) and 8,250 sec-



onds of slack time in Process 3 (where nuts and raisins are mixed and added to the chocolate).

**Constraint Summary**

The last section of the output is the CONSTRAINT SUMMARY. This output (Figure 1.11) gives the value of the objective function, the value of each constraint, and the dual activities.

The ACTIVITY column gives the value of the right-hand side of each equation when the problem is solved using the information given in the previous table.

The DUAL ACTIVITY column reveals that each second in Process 1 (mixing-cooking) is worth approximately \$.013, and each second in Process 4 (Packaging), is worth approximately \$.005. These figures (called *shadow prices*) can be used to decide whether the total available time for Processes 1 and 4 should be increased. If a second can be added to the total production time in Processes 1 for less than \$.013, it would be profitable to do so. The dual activities for Process 1 and 4 are zero since adding time to those processes does not increase profits. Keep in mind that the dual activity gives the marginal improvement to the objective and that adding time to Process 1 changes the original problem and solution.

The LP Procedure						
Constraint Summary						
Row	Constraint Name	Type	S/S Col	Rhs	Activity	Dual Activity
1	object	OBJECTIVE	.	0	475	.
2	process1	LE	3	27000	27000	0.012963
3	process2	LE	4	27000	16875	0
4	process3	LE	5	27000	18750	0
5	process4	LE	6	27000	27000	0.0046296

**Figure 1.11.** CONSTRAINT SUMMARY

For a complete description of the output from PROC LP, see the chapter on teh LP procedure.

**Sparse Format**

Typically, mathematical programming models are sparse. That is, few of the coefficients in the constraint matrix are nonzero. The dense problem format shown in the last section would be an inefficient way to represent sparse models. The LP procedure also accepts data in a sparse input format. Only the nonzero coefficients need be specified. It is consistent with the standard MPS sparse format and much more flexible so that models using the MPS format can be easily converted to the LP format. The appendix at the end of this book describes a SAS macro for conversion.

Although the factory example of the last section is not sparse, illustrated here is an example of the sparse input format. The sparse data set has four variables: a row type identifying variable, a row name variable, a column name variable, and a coefficient variable.

```

data factory;
  format _type_ $8. _row_ $16. _col_ $16.;
  input _type_ $ _row_ $ _col_ $ _coef_ ;
  datalines;
max      object      .      .
.        object      chocolate  .25
.        object      toffee     .75
le       process1    .        .
.        process1    chocolate  15
.        process1    toffee     40
.        process1    _RHS_      27000
le       process2    .        .
.        process2    toffee     56.25
.        process2    _RHS_      27000
le       process3    .        .
.        process3    chocolate  18.75
.        process3    _RHS_      27000
le       process4    .        .
.        process4    chocolate  12
.        process4    toffee     50
.        process4    _RHS_      27000
;

```

To solve this problem using the Interior Point algorithm of PROC NETFLOW, specify

```
proc netflow sparsesecondata arcdata=factory condata=factory;
```

However, this example will be solved by the LP procedure.

Notice that the `_type_` variable contains keywords as for the dense format, the `_row_` variable contains the row names in the model, the `_col_` variable contains the column names in the model, and the `_coef_` variable contains the coefficients for that particular row and column. Since the row and column names are the values of variables in a SAS data set, they are not limited to 8 characters. This feature, as well as the order independence of the format, simplify matrix generation.

The `SPARSEDATA` option in the PROC LP statement tells the LP procedure that the model in the problem data set is in the sparse format. This example also illustrates how the solution of the linear program is saved in two output data sets: the primal data set and the dual data set.

```

proc lp
  data=factory sparsedata
  primalout=primal dualout=dual;
run;

```

The primal data set (Figure 1.12) contains the information that is displayed in the `VARIABLE SUMMARY` plus additional information about the bounds on the variables.

Obs	_OBJ_ID_	_RHS_ID_	_VAR_	_TYPE_	_STATUS_
1	object	_RHS_	chocolate	NON-NEG	_BASIC_
2	object	_RHS_	toffee	NON-NEG	_BASIC_
3	object	_RHS_	process1	SLACK	
4	object	_RHS_	process2	SLACK	_BASIC_
5	object	_RHS_	process3	SLACK	_BASIC_
6	object	_RHS_	process4	SLACK	
7	object	_RHS_	PHASE_1_OBJECTIV	OBJECT	_DEGEN_
8	object	_RHS_	object	OBJECT	_BASIC_

Obs	_LBOUND_	_VALUE_	_UBOUND_	_PRICE_	_R_COST_
1	0	1000	1.7977E308	0.25	-0.000000
2	0	300	1.7977E308	0.75	0.000000
3	0	0	1.7977E308	0.00	-0.012963
4	0	10125	1.7977E308	0.00	0.000000
5	0	8250	1.7977E308	0.00	0.000000
6	0	0	1.7977E308	0.00	-0.004630
7	0	0	0	0.00	0.000000
8	0	475	1.7977E308	0.00	0.000000

**Figure 1.12.** Primal Data Set


The dual data set (Figure 1.13) contains the information that is displayed in the CONSTRAINT SUMMARY plus additional information about bounds on the rows.

Obs	_OBJ_ID_	_RHS_ID_	_ROW_ID_	_TYPE_	_RHS_	_L_RHS_	_VALUE_	_U_RHS_	_DUAL_
1	object	_RHS_	object	OBJECT	475	475	475	475	.
2	object	_RHS_	process1	LE	27000	-1.7977E308	27000	27000	0.012963
3	object	_RHS_	process2	LE	27000	-1.7977E308	16875	27000	0.000000
4	object	_RHS_	process3	LE	27000	-1.7977E308	18750	27000	0.000000
5	object	_RHS_	process4	LE	27000	-1.7977E308	27000	27000	0.004630

**Figure 1.13.** Dual Data Set

**Network Format**

Network flow problems can be described by specifying the nodes in the network and their supplies and demands, and the arcs in the network and their costs and capacities and lower flow bounds. Consider the simple transshipment problem in Figure 1.14 as an illustration.



**Figure 1.14.** Transshipment Problem

Suppose the candy manufacturing company has two factories, two warehouses, and three customers for chocolate. The two factories each have a production capacity of 500 pounds per day. The three customers have demands of 100, 200, and 50 pounds per day, respectively.

The following data set describes the supplies (positive values for the `supdem` variable) and the demands (negative values for the `supdem` variable) for each of the customers and factories.

```

data nodes;
  format node $16. ;
  input node $ supdem;
  datalines;
customer_1  -100
customer_2  -200
customer_3   -50
factory_1    500
factory_2    500
;
  
```

Suppose that there are two warehouses that are used to store the chocolate before shipment to the customers and that there are different costs for shipping between each factory, warehouse, and customer.

What is the minimum cost routing for supplying the customers?

Arcs are described in another data set. Each observation defines a new arc in the network and gives data about the arc. For example, there is an arc between the node

factory\_1 and the node warehouse\_1. Each unit of flow on that arc costs 10. Although this example does not include it, lower and upper bounds on the flow across that arc can be listed here.

```
data network;
  format from $16. to $16.;
  input from $ to $ cost ;
  datalines;
factory_1      warehouse_1  10
factory_2      warehouse_1   5
factory_1      warehouse_2   7
factory_2      warehouse_2   9
warehouse_1    customer_1   3
warehouse_1    customer_2   4
warehouse_1    customer_3   4
warehouse_2    customer_1   5
warehouse_2    customer_2   5
warehouse_2    customer_3   6
  ;
```

Use PROC NETFLOW to find the minimum cost routing. This procedure takes the model as defined in the network and nodes data sets and finds the minimum cost flow.

```
proc netflow arcout=arc_sav
             arcdata=network nodedata=nodes;
  node node;      /* node data set information */
  supdem supdem;

  tail from;      /* arc data set information */
  head to;
  cost cost;

  run;

proc print;
  var from to cost _capac_ _lo_ _supply_ _demand_
      _flow_ _fcost_ _rcost_;
  sum _fcost_;
run;
```

PROC NETFLOW produces the following on the SAS log:

```
NOTE: Number of nodes= 7 .
NOTE: Number of supply nodes= 2 .
NOTE: Number of demand nodes= 3 .
NOTE: Total supply= 1000 , total demand= 350 .
NOTE: Number of arcs= 10 .
NOTE: Number of iterations performed (neglecting any constraints)= 7 .
NOTE: Of these, 2 were degenerate.
NOTE: Optimum (neglecting any constraints) found.
NOTE: Minimal total cost= 3050 .
NOTE: The data set WORK.ARC_SAV has 10 observations and 13 variables.
```

The solution (Figure 1.15) saved in the `arc_sav` data set shows the optimal amount of chocolate to send across each arc (the amount to ship from each factory to each warehouse and from each warehouse to each customer) in the network per day.

f		c		-		S		D		-		F		R	
O		o		A		U		E		-		C		C	
b		s		P		P		M		L		O		O	
o		t		A		L		A		O		S		S	
s		o		C		O		D		W		T		T	
m		t		-		-		-		-		-		-	
1	warehouse_1	customer_1	3	99999999	0	.	100	100	300	.					
2	warehouse_2	customer_1	5	99999999	0	.	100	0	0	4					
3	warehouse_1	customer_2	4	99999999	0	.	200	200	800	.					
4	warehouse_2	customer_2	5	99999999	0	.	200	0	0	3					
5	warehouse_1	customer_3	4	99999999	0	.	50	50	200	.					
6	warehouse_2	customer_3	6	99999999	0	.	50	0	0	4					
7	factory_1	warehouse_1	10	99999999	0	500	.	0	0	5					
8	factory_2	warehouse_1	5	99999999	0	500	.	350	1750	.					
9	factory_1	warehouse_2	7	99999999	0	500	.	0	0	.					
10	factory_2	warehouse_2	9	99999999	0	500	.	0	0	2					
										====					
										3050					

Figure 1.15. ARCOUT Data Set

Notice which arcs have positive flow (`_FLOW_` is greater than 0). These arcs indicate the amount of chocolate that should be sent from `factory_2` to `warehouse_1` and from there on to the three customers. The model indicates no production at `factory_1` and no use of `warehouse_2`.




Figure 1.16. Optimal Solution for the Transshipment Problem

## PROC ASSIGN and PROC TRANS Data Formats

The transportation and assignment models are described in rectangular data sets. Suppose that instead of sending chocolate from factories to warehouses and then to the customers, chocolate is sent directly from the factories to the customers.

Finding the minimum cost routing could be done using the NETFLOW procedure. However, since the network represents a transportation problem, the data for the problem can be represented more simply.

```
data transprt;
  input source $ supply cust_1 cust_2 cust_3 ;
  datalines;
demand      .    100    200    50
factory1    500    10     9     7
factory2    500     9    10     8
  ;
```

This data set shows the source names as the values for the `SOURCE` variable, the supply at each source node as the values for the `SUPPLY` variable, and the unit shipping cost for source to sink as the values for the sink variables `CUST_1` to `CUST_3`. Notice that the first record contains the demands at each of the sink nodes.

The TRANS procedure finds the minimum cost routing. It solves the problem and saves the solution in an output data set.

```
proc trans
  nothrunet data=transprt out=transout;
  supply supply;
  id source;

proc print;run;
```

The optimum solution total (3050) is reported on the SAS log. The entire solution (Figure 1.17) shows the amount of chocolate to ship from each factory to each customer, per day.

The resulting data set called `OUT` contains the variables listed in the `VAR` and `ID` statements, and a new variable called `_DUAL_`. This variable `_DUAL_` contains the marginal costs of increasing the supply at each origin point. The last observation in the `OUT` data set has the marginal costs of increasing the demand at each destination point. These variables are called dual variables.

Obs	source	supply	cust_1	cust_2	cust_3	_DUAL_
1	_DEMAND_	.	100	200	50	.
2	factory1	500	0	200	50	0
3	factory2	500	100	0	0	0
4	_DUAL_	.	9	9	7	.

Figure 1.17. PROC TRANS Solution

## Model Building

It is often desirable to keep the data separate from the structure of the model. This is useful for large models with numerous identifiable components. The data are best organized in rectangular tables that can be easily examined and modified. Then, before the problem is solved, the model is built using the stored data. This process of model building is known as *matrix generation*. In conjunction with the sparse format, the SAS DATA step provides a good matrix generation language.

For example, consider the candy manufacturing example introduced previously. Suppose that for the user interface it is more convenient to have the data so that each record describes the information related to each product (namely, the contribution to the objective function and the unit amount needed for each process) needed by each product. A DATA step for saving these data might look like this:

```
data manfg;
  input product $16. object process1 - process4 ;
  datalines;
chocolate      .25    15  0.00 18.75    12
toffee         .75    40 56.25  0.00    50
licorice       1.00   29 30.00 20.00    20
jelly_beans   .85    10  0.00 30.00    10
_RHS_         .    27000 27000 27000 27000
  ;
```

Notice that there is a special record at the end having product `_RHS_`. This record gives the amounts of time available for each of the processes. This information could have been stored in another data set. The next example illustrates a model where the data are stored in separate data sets.

Building the model involves adding the data to the structure. There are as many ways to do this as there are programmers and problems. The following DATA step shows one way to take the candy data and build a sparse format model to solve the product mix problem.

```
data model;
  array process object process1-process4;
  format _type_ $8. _row_ $16. _col_ $16. ;
  keep _type_ _row_ _col_ _coef_;

  set manfg;          /* read the manufacturing data */

  /* build the object function */

  if _n_=1 then do;
    _type_='max'; _row_='object'; _col_=' '; _coef_=.; output;
  end;

  /* build the constraints */

  do over process;
```



```

if _i_>1 then do;
  _type_='le'; _row_='process' || put(_i_,1.);
end;
else          _row_='object';
_col_=product; _coef_=process;
output;
end;
run;

```

The sparse format data set is shown in Figure 1.18.

Obs	_type_	_row_	_col_	_coef_
1	max	object		.
2	max	object	chocolate	0.25
3	le	process1	chocolate	15.00
4	le	process2	chocolate	0.00
5	le	process3	chocolate	18.75
6	le	process4	chocolate	12.00
7		object	toffee	0.75
8	le	process1	toffee	40.00
9	le	process2	toffee	56.25
10	le	process3	toffee	0.00
11	le	process4	toffee	50.00
12		object	licorice	1.00
13	le	process1	licorice	29.00
14	le	process2	licorice	30.00
15	le	process3	licorice	20.00
16	le	process4	licorice	20.00
17		object	jelly_beans	0.85
18	le	process1	jelly_beans	10.00
19	le	process2	jelly_beans	0.00
20	le	process3	jelly_beans	30.00
21	le	process4	jelly_beans	10.00
22		object	_RHS_	.
23	le	process1	_RHS_	27000.00
24	le	process2	_RHS_	27000.00
25	le	process3	_RHS_	27000.00
26	le	process4	_RHS_	27000.00

**Figure 1.18.** Sparse Data Format

The model data set looks a little different from the sparse representation of the candy model shown earlier. It not only includes the additional products, licorice and jelly\_beans, but it also defines the model in a different order. Since the sparse format is robust, the model can be generated in ways that are convenient for the data step program.

If the problem had more products, increase the size of the `manfg` data set to include the new product data. Also, if the problem had more than 4 processes, add the new process variables to the `manfg` data set and increase the size of the `process` array in the model data set. With these two simple changes and additional data, a product mix problem having hundreds of processes and products can be solved.

## Matrix Generation

It is desirable to keep data in separate tables, then automate model building and reporting. This example illustrates a problem that has elements of a product mix problem and a blending problem. Suppose four kinds of ties are made; all silk, all polyester, a 50-50 polyester cotton blend, and a 30-70 cotton polyester blend.

The data includes cost and supplies of raw material; selling price, minimum contract sales, and maximum demand of the finished products; and the proportions of raw materials that go into each product. The product mix that maximizes profit is to be found.

The data are saved in three SAS data sets. The program that follows demonstrates one way for these data to be saved. Alternatively, the full screen editor PROC FSEDIT can be used to store and edit these data can be used.

```

data material;
  input descpt $char20. cost supply;
datalines;
silk_material          .21  25.8
polyester_material     .6   22.0
cotton_material        .9   13.6
;

data tie;
  input descpt $char20. price contract demand;
datalines;
all_silk                6.70   6.0   7.00
all_polyester           3.55  10.0  14.00
poly_cotton_blend      4.31  13.0  16.00
cotton_poly_blend      4.81   6.0   8.50
;

data manfg;
  input descpt $char20. silk poly cotton;
datalines;
all_silk                100    0    0
all_polyester           0   100   0
poly_cotton_blend       0    50   50
cotton_poly_blend       0    30   70
;

```

The program that follows takes the raw data from the three data sets and builds a linear program model in the data set called `model`. Although it is designed for the three-resource, four-product problem described here, it can be easily extended to include more resources and products. The model building DATA step remains essentially the same, all that changes are the dimensions of loops and arrays. Of course the data tables must increase to accommodate the new data.

```

data model;
  array raw_mat {3} $ 20 ;
  array raw_comp {3} silk poly cotton;
  length _type_ $ 8 _col_ $ 20 _row_ $ 20 _coef_ 8 ;
  keep _type_ _col_ _row_ _coef_ ;

  /* define the objective, lower, and upper bound rows */

  _row_='profit'; _type_='max'; output;
  _row_='lower'; _type_='lowerbd'; output;
  _row_='upper'; _type_='upperbd'; output;
  _type_=' ';

  /* the object and upper rows for the raw materials */

do i=1 to 3;
  set material;
  raw_mat[i]=descpt; _col_=descpt;
  _row_='profit'; _coef_=-cost; output;
  _row_='upper'; _coef_=supply; output;
end;

  /* the object, upper, and lower rows for the products */

do i=1 to 4;
  set tie;
  _col_=descpt;
  _row_='profit'; _coef_=price; output;
  _row_='lower'; _coef_=contract; output;
  _row_='upper'; _coef_=demand; output;
end;

  /* the coefficient matrix for manufacturing */

  _type_='eq';
do i=1 to 4; /* loop for each raw material */
  set manfg;
  do j=1 to 3; /* loop for each product */

    _col_=descpt; /* % of material in product */
    _row_ = raw_mat[j];
    _coef_ = raw_comp[j]/100;
    output;

    _col_ = raw_mat[j]; _coef_ = -1;
    output;

  /* the right-hand-side */

  if i=1 then do;
    _col_='_RHS_';
    _coef_=0;
    output;
  end;
end;

```

```

        end;
        _type_='  ';
    end;
stop;
run;

```

The model is solved using PROC LP which saves the solution in the PRIMALOUT data set named solution. PROC PRINT displays the solution, which is shown in Figure 1.19.

```

proc lp sparsedata primalout=solution;

proc print ;
    id _var_;
    var _lbound_--_r_cost_;
run;

```

The results are shown in Figure 1.19.

_VAR_	_LBOUND_	_VALUE_	_UBOUND_	_PRICE_	_R_COST_
all_polyester	10	11.800	14.0	3.55	0.000
all_silk	6	7.000	7.0	6.70	6.490
cotton_material	0	13.600	13.6	-0.90	4.170
cotton_poly_blend	6	8.500	8.5	4.81	0.196
polyester_material	0	22.000	22.0	-0.60	2.950
poly_cotton_blend	13	15.300	16.0	4.31	0.000
silk_material	0	7.000	25.8	-0.21	0.000
PHASE_1_OBJECTIVE	0	0.000	0.0	0.00	0.000
profit	0	168.708	1.7977E308	0.00	0.000

**Figure 1.19.** The Solution Data Set

The solution shows that 11.8 units of polyester ties, 7 units of silk ties, 8.5 units of the cotton-polyester blend, and 15.3 units of the polyester-cotton blend should be produced. It also shows the amounts of raw materials that go into this product mix to generate a total profit of 168.708.

## Exploiting Model Structure

Another example helps to illustrate how the model can be simplified by exploiting the structure in the model and by using the NETFLOW procedure.

Recall the chocolate transshipment problem discussed earlier. The solution required no production at factory\_1 and no storage at warehouse\_2. Suppose this solution, although optimal, is unacceptable. An additional constraint requiring the production at the two factories to be balanced is required. Now, the production at the two factories can differ by, at most, 100 units. Such a constraint might look like

$$-100 \leq (\text{factory\_1\_warehouse\_1} + \text{factory\_1\_warehouse\_2} - \text{factory\_2\_warehouse\_1} - \text{factory\_2\_warehouse\_2}) \leq 100$$

The network and supply and demand information is saved in the two data sets.

```

data network;
  format from $16. to $16.;
  input from $ to $ cost ;
  datalines;
factory_1 warehouse_1 10
factory_2 warehouse_1 5
factory_1 warehouse_2 7
factory_2 warehouse_2 9
warehouse_1 customer_1 3
warehouse_1 customer_2 4
warehouse_1 customer_3 4
warehouse_2 customer_1 5
warehouse_2 customer_2 5
warehouse_2 customer_3 6
;

data nodes;
  format node $16. ;
  input node $ supdem;
  datalines;
customer_1 -100
customer_2 -200
customer_3 -50
factory_1 500
factory_2 500
;

```

The factory balancing constraint is not a part of the network. It is represented in the sparse format in a data set for side constraints.

```

data side_con;
  format _type_ $8. _row_ $16. _col_ $21. _coef_ 8.;
  input _type_ _row_ _col_ _coef_ ;
  datalines;
eq balance . .
. balance factory_1_warehouse_1 1
. balance factory_1_warehouse_2 1
. balance factory_2_warehouse_1 -1
. balance factory_2_warehouse_2 -1
. balance diff -1
lo lowerbd diff -100
up upperbd diff 100
;

```

It contains an equality constraint that sets the value of DIFF to be the amount that factory 1 production exceeds factory 2 production. It also contains implicit bounds on the DIFF variable. Note that the DIFF variable is a nonarc variable.



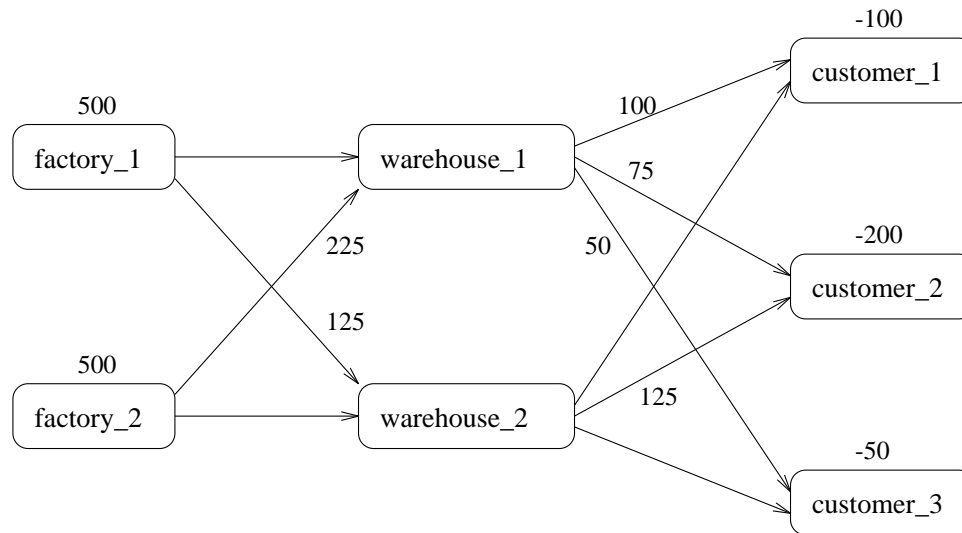


Figure 1.21. Constrained Optimum for the Transshipment Problem

---

## Report Writing

The reporting of the solution is also an important aspect of modeling. Since the optimization procedures save the solution in one or more SAS data sets, report writing can be written using any of the tools in the SAS language.

---

## The DATA step

Use of the DATA step and PROC PRINT is the most general way to produce reports. For example, a table showing the revenue generated from the production and a table of the cost of material can be produced with the following program.

```

data product(keep= _var_ _value_ _price_ revenue)
  material(keep= _var_ _value_ _price_ cost);
set solution;
if _price_>0 then do;
  revenue=_price*_value_; output product;
end;
else if _price_<0 then do;
  _price_=-_price_;
  cost = _price*_value_; output material;
end;
run;

/* display the product report */

```

```

proc print data=product;
  id _var_;
  var _value_ _price_ revenue ;
  sum revenue;
  title 'Revenue Generated from Tie Sales';
run;

/* display the materials report */

proc print data=material;
  id _var_;
  var _value_ _price_ cost;
  sum cost;
  title 'Cost of Raw Materials';
run;

```

This DATA step reads the solution data set saved by PROC LP and segregates the records based on whether they correspond to materials or products, namely whether the contribution to profit is positive or negative. Each of these is then displayed to produce Figure 1.22.

Revenue Generated from Tie Sales			
_VAR_	_VALUE_	_PRICE_	revenue
all_polyester	11.8	3.55	41.890
all_silk	7.0	6.70	46.900
cotton_poly_blend	8.5	4.81	40.885
poly_cotton_blend	15.3	4.31	65.943
			=====
			195.618

Cost of Raw Materials			
_VAR_	_VALUE_	_PRICE_	cost
cotton_material	13.6	0.90	12.24
polyester_material	22.0	0.60	13.20
silk_material	7.0	0.21	1.47
			=====
			26.91

Figure 1.22. Tie problem: Revenues and Costs

## Other Reporting Procedures


The GCHART procedure can be a useful tool for displaying the solution to mathematical programming models. The `con_solv` data set that contains the solution to the balanced transshipment problem can be effectively displayed using PROC GCHART. In Figure 1.23, the amount that is shipped from each factory and warehouse can be seen by submitting the following.



```

title;
proc gchart data=con_sav;
  hbar from / sumvar=_flow_;
run;

```




**Figure 1.23.** Tie problem: Through-puts

The horizontal bar chart is just one way of displaying the solution to a mathematical program. The solution to the Tie Product Mix problem that was solved using PROC LP can also be illustrated using PROC GCHART. Here, a pie chart shows the relative contribution of each product to total revenues.

```

proc gchart data=product;
  pie _var_ / sumvar=revenue;
  title 'Projected Tie Sales Revenue';
run;

```



**Figure 1.24.** Tie Problem: Projected Tie Sales Revenue

The TABULATE procedure is another procedure that can help automate solution reporting. There are several examples in the chapter on the LP procedure that illustrate its use.

---

## Decision Support Systems

The close relationship between a SAS data set and the representation of the mathematical model make it easy to build decision support systems.

---

### The Full-Screen Interface

The ability to manipulate data using the full screen tools in the SAS language further enhances the decision support capabilities. The several data set pieces that are components of a decision support model can be edited using the full screen editing procedures FSEDIT and FSPRINT. The screen control language SCL directs data editing, model building, and solution reporting through its menuing capabilities.

The compatibility of each of these pieces in the SAS System make construction of a fullscreen decision support system based on mathematical programming an easy task.

---

## Communicating with the Optimization Procedures

The optimization procedures communicate with any decision support system through the various problem and solution data sets. However, there is a need for the system to have a more intimate knowledge of the status of the optimization procedures. This is achieved through the use of macro variables defined by each of the optimization procedures.

---

## References

Rosenbrock, H.H. (1960), “An Automatic Method for Finding the Greatest or Least Value of a Function,” *Comput. J.*, **3**, 175–184.

The correct bibliographic citation for this manual is as follows: SAS Institute Inc., *SAS/OR<sup>®</sup> User's Guide: Mathematical Programming, Version 8*, Cary, NC: SAS Institute Inc., 1999. 566 pp.

**SAS/OR<sup>®</sup> User's Guide: Mathematical Programming, Version 8**

Copyright © 1999 SAS Institute Inc., Cary, NC, USA.

ISBN 1-58025-491-8

All rights reserved. Printed in the United States of America. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, by any form or by any means, electronic, mechanical, photocopying, or otherwise, without the prior written permission of the publisher, SAS Institute, Inc.

**U.S. Government Restricted Rights Notice.** Use, duplication, or disclosure of the software by the government is subject to restrictions as set forth in FAR 52.227-19 Commercial Computer Software-Restricted Rights (June 1987).

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

1st printing, October 1999

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

IBM<sup>®</sup>, ACF/VTAM<sup>®</sup>, AIX<sup>®</sup>, APPN<sup>®</sup>, MVS/ESA<sup>®</sup>, OS/2<sup>®</sup>, OS/390<sup>®</sup>, VM/ESA<sup>®</sup>, and VTAM<sup>®</sup> are registered trademarks or trademarks of International Business Machines Corporation. <sup>®</sup> indicates USA registration.

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

The Institute is a private company devoted to the support and further development of its software and related services.