Chapter 3 The LP Procedure

Chapter Table of Contents

OVERVIEW
GETTING STARTED 57 An Introductory Example 58 An Integer Programming Example 62 An MPS Example 64
An MPS Format Conversion Example
SYNTAX
Functional Summary
PROC LP Statement
COEF Statement
COL Statement
ID Statement
IPIVOT Statement 82
PIVOT Statement
PRINT Statement
QUIT Statement
RANGE Statement
RESET Statement
RHS Statement
RHSSEN Statement
ROW Statement
RUN Statement
SHOW Statement
TYPE Statement 88
VAR Statement
Missing Volues 01
MISSING Values 91 Sparse Data Input Format 91
Converting MDS Formet
The Deduced Costs Duck Activities and Comment Tableco
Magna Mariahla ODLD
Macro variable _OKLP
Pricing
Scanng
Preprocessing

Integer Programming
Sensitivity Analysis
Range Analysis
Parametric Programming
Interactive Facilities
Memory Management
Output Data Sets
Input Data Sets
Displayed Output
ODS Table and Variable Names
EXAMPLES
Example 3.1 An Oil Blending Problem
Example 3.2 A Sparse View of the Oil Blending Problem
Example 3.3 Analyzing the Sensitivity of the Solution to Changes in the
Objective Coefficients
Example 3.4 Additional Analysis of the Sensitivity of the Solution to
Changes in the Objective Coefficients
Example 3.5 Price Parametric Programming for the Oil Blending Problem 134
Example 3.6 Special Ordered Sets and the Oil Blending Problem
Example 3.7 Goal-Programming a Product Mix Problem
Example 3.8 A Simple Integer Program
Example 3.9 An Infeasible Problem
Example 3.10 Restarting an Integer Program
Example 3.11 Alternative Search of the Branch and Bound Tree
Example 3.12 An Assignment Problem
Example 3.13 A Scheduling Problem
Example 3.14 A Multicommodity Transshipment Problem with Fixed Charges 177
REFERENCES

Chapter 3 The LP Procedure

Overview

The LP procedure solves linear programs, integer programs, and mixed-integer programs. It also performs parametric programming, range analysis, and reports on solution sensitivity to changes in the right-hand-side constants and price coefficients.

The LP procedure provides various control options and solution strategies. It also provides the functionality to produce various kinds of intermediate and final solution information. The procedure's interactive features enable you to take control of the problem solving process. During linear or integer iterations, for example, you can stop the procedure at intermediate stages and examine current results. If necessary, you can change options or strategies and resume the execution of the procedure.

The LP procedure is used to optimize a linear function subject to linear and integer constraints. Specifically, the LP procedure solves the general mixed-integer program of the form

$$\min c^T x$$

subject to

$$Ax \ge =, \le b$$

 $\ell \le x \le u$
 $x_{i \in S}$ is integer

where

A is an $m \ge n$ matrix of technological coefficients.

b is an $m \ge 1$ matrix of right-hand-side (RHS) constants.

c is an $n \ge 1$ matrix of objective function coefficients.

x is an $n \ge 1$ matrix of structural variables.

l is an $n \ge 1$ matrix of lower bounds on x.

u is an $n \ge 1$ matrix of upper bounds on x.

S is a subset of the set of indices $\{1, \ldots, n\}$.

56 • Chapter 3. The LP Procedure

Linear programs (when S is empty) are denoted by (lp). For these problems, the procedure employs the two-phase revised simplex method, which uses the Bartels-Golub update of the LU decomposed basis matrix to pivot between feasible solutions (Bartels 1971). In phase 1, PROC LP finds a basic feasible solution to (lp), while in phase 2, PROC LP finds an optimal solution, x^{opt} . The procedure implicitly handles unrestricted variables, lower-bounded variables, upper-bounded variables, and ranges on constraints. When no explicit lower bounds are specified, PROC LP assumes that all variables are bounded below by zero.

When a variable is specified as an integer variable, S has at least one element. Then, the procedure uses the branch and bound technique for optimization.

The relaxed problem (the problem with no integer constraints) is solved initially using the primal algorithm described previously. Constraints are added in defining the subsequent descendent problems in the branch and bound tree. These problems are then solved using the dual simplex algorithm. Dual pivots are referred to as phase 3 pivots.

The preprocessing option enables the procedure to identify redundant and infeasible constraints, fix variables, and reduce the feasible region before solving a problem. For linear programs, the option often can reduce the number of constraints and variables, leading to a quicker elapsed solution time and improved reliability. For integer programs, it often reduces the gap between an integer program and its relaxed linear program, which will likely lead to a reduced branch and bound tree and a quicker CPU time. In general, it provides users an alternative to solving large, complicated operations research problems.

The LP procedure can also analyze the sensitivity of the solution x^{opt} to changes in both the objective function and the right-hand-side constants. There are three techniques available for this analysis: sensitivity analysis, parametric programming, and range analysis. Sensitivity analysis enables you to examine the size of a perturbation to the right-hand-side or objective vector by an arbitrary change vector for which the basis of the current optimal solution remains optimal.

Parametric programming, on the other hand, enables you to specify the size of the perturbation beforehand and examine how the optimal solution changes as the desired perturbation is realized. With this technique, the procedure pivots to maintain optimality as the right-hand-side or objective vector is perturbed beyond the range for which the current solution is optimal. Range analysis is used to examine the range of each right-hand-side value or objective coefficient for which the basis of the current optimal solution remains optimal.

The LP procedure can also save both primal and dual solutions, the current tableau, and the branch and bound tree in SAS data sets. This enables you to generate solution reports and perform additional analyses with the SAS System. Although PROC LP reports solutions, this feature is particularly useful for reporting solutions in formats tailored to your specific needs. Saving computational results in a data set also enables you to continue executing a problem not solved because of insufficient time or other computational problems.

The LP procedure uses the Output Delivery System (ODS), a SAS subsystem that provides capabilities for displaying and controlling the output from SAS procedures. ODS enables you to modify the headers, column names, data formats, and layouts of the output tables in PROC LP.

There are no restrictions on the problem size in the LP procedure. The number of constraints and variables in a problem that PROC LP can solve depends on the host platform, the available memory, and the available disk space for utility data sets.

Getting Started

PROC LP expects the definition of one or more linear, integer, or mixed-integer programs in an input data set. There are two formats, a dense format and a sparse format, for this data set.

In the dense format, a model is expressed in a similar way as it is formulated. Each SAS variable corresponds to a model's column, and each SAS observation corresponds to a model's row. A SAS variable in the input data set is one of the following:

- a type variable
- an id variable
- a structural variable
- a right-hand-side variable
- a right-hand-side sensitivity analysis variable or
- a range variable

The type variable tells PROC LP how to interpret the observation as a part of the mathematical programming problem. It identifies and classifies objectives, constraints, and the rows that contain information of variables like types, bounds, and so on. PROC LP recognizes the following keywords as values for the type variable: *MIN*, *MAX*, *EQ*, *LE*, *GE*, *SOSEQ*, *SOSLE*, *UNRSTRCT*, *LOWERBD*, *UPPERBD*, *FIXED*, *INTEGER*, *BINARY*, *BASIC*, *PRICESEN* and *FREE*. The values of the id variable are the names of the rows in the model. The other variables identify and classify the columns with numerical values.

The sparse format to PROC LP is designed to enable you to specify only the nonzero coefficients in the description of linear programs, integer programs, and mixed-integer programs. The SAS data set that describes the sparse model must contain at least four SAS variables:

- a type variable
- a column variable
- a row variable and
- a coefficient variable

Each observation in the data set associates a type with a row or a column, or defines a coefficient or a numerical value in the model, or both. In addition to the keywords

in the dense format, PROC LP also recognizes the keywords *RHS*, *RHSSEN* and *RANGE* as values of the type variable. The values of the row and column variables are the names of the rows and columns in the model. The values of the coefficient variables give the coefficients or other numerical data. The SAS data set can contain multiple pairs of row and coefficient variables. In this way, more information about the model can be specified in each observation in the data set. See the "Sparse Data Input Format" section on page 91 for further discussion.

With both the dense and sparse formats for model specification, the observation order is not important. This feature is particularly useful when using the sparse model input.

An Introductory Example

A simple blending problem illustrates the dense and sparse input formats and the use of PROC LP. A step in refining crude oil into finished oil products involves a distillation process that splits crude into various streams. Suppose there are three types of crude available: Arabian light, Arabian heavy, and Brega. These types of crude are distilled into light naphtha, intermediate naphtha, and heating oil. These in turn are blended into jet fuel using one of two recipes. What amounts of the three crudes maximize the profit from producing jet fuel? A formulation to answer this question is as follows:

```
\max-175a\_light-165a\_heavy-205brega+300jet\_1+300jet\_2
```

s.t.
$$.035a_light + .03a_heavy + .045brega = naphthal$$

 $.1a_light + .075a_heavy + .135brega = naphthai$
 $.39a_light + .3a_heavy + .43brega = heatingo$
 $.3naphthai + .7heatingo = jet_1$
 $.2naphthal + .8heatingo = jet_2$
 $a_light \leq 110$
 $a_heavy \leq 165$
 $brega \leq 80$
 $a_light, a_heavy, brega, naphthai,$
 $naphthal, heatingos, jet_1, jet_2 \geq 0$

The following data set gives the representation of this formulation. Notice that the variable names are the structural variables, the rows are the constraints, and the coefficients are given as the values for the structural variables.

```
data;
input _id_ $14.
    a_light a_heavy brega naphthal naphthai
    heatingo jet_1 jet_2
    _type_ $ _rhs_;
```

datalines;										
profit	-175	-165	-205	0	0	0	300	300	max	•
naphtha_l_conv	.035	.030	.045	-1	0	0	0	0	eq	0
naphtha_i_conv	.100	.075	.135	0	-1	0	0	0	eq	0
heating_o_conv	.390	.300	.430	0	0	-1	0	0	eq	0
recipe_1	0	0	0	0	.3	.7	-1	0	eq	0
recipe_2	0	0	0	.2	0	.8	0	-1	eq	0
available	110	165	80	•	•	•	•	•	upperbd	•
;										

The same model can be specified in the sparse format, as follows. This format enables you to omit the zero coefficients.

data;			
input	_type_ \$ @10	_col_ \$13. @24 _row_	\$16coef_;
datal	ines;		
max	•	profit	•
eq	•	napha_1_conv	•
eq	•	napha_i_conv	•
eq	•	heating_oil_conv	•
eq	•	recipe_1	•
eq	•	recipe_2	•
upperbd	•	available	•
•	a_light	profit	-175
•	a_light	napha_l_conv	.035
•	a_light	napha_i_conv	.100
•	a_light	heating_oil_conv	.390
•	a_light	available	110
•	a_heavy	profit	-165
•	a_heavy	napha_l_conv	.030
•	a_heavy	napha_i_conv	.075
•	a_heavy	heating_oil_conv	.300
•	a_heavy	available	165
•	brega	profit	-205
•	brega	napha_l_conv	.045
•	brega	napha_i_conv	.135
•	brega	heating_oil_conv	.430
•	brega	available	80
•	naphthal	napha_l_conv	-1
•	naphthal	recipe_2	.2
•	naphthai	napha_i_conv	-1
•	naphthai	recipe_1	.3
•	heatingo	heating_oil_conv	-1
•	heatingo	recipe_1	.7
•	heatingo	recipe_2	.8
•	jet_1	profit	300
•	jet_1	recipe_1	-1
•	jet_2	profit	300
•	jet_2	recipe_2	-1
•	_rhs_	recipe_1	0
;			

Because the input order of the model into PROC LP is unimportant, this model can be specified in sparse input in arbitrary row order. Example 3.2 in the "Examples" section demonstrates this.

The dense and sparse forms of model input give you flexibility to generate models using the SAS language. The dense form of the model is solved with the statements

proc lp;
run;

The sparse form is solved with the statements:

```
proc lp sparsedata;
run;
```

Example 3.1 and Example 3.2 in the "Examples" section continue with this problem.

Problem Input

In the default, PROC LP uses the most recently created SAS data set as the problem input data set. However, if you want to input the problem from a specific SAS data set, then use the DATA= option. For example, if the previous dense form data set has a name, DENSE, the PROC LP statements can be written as

proc lp data=dense; run;

Problem Definition Statements

In the previous dense form data set, the _ID_, _TYPE_, and _RHS_ variables are special variables in PROC LP. They stand for id variable, type variable, and right-hand-side variable. If you replace those variable names with names like ROWNAME, TYPE, and RHS, you need the problem definition statements (ID, TYPE and RHS) in PROC LP:

```
proc lp;
    id rowname;
    type type;
    rhs rhs;
run;
```

Other special variables for the dense format are _RHSSEN_ and _RANGE_, which identify the vectors for the right-hand-side sensitivity and range analyses. The corresponding statements are the RHSSEN and RANGE statements. (Notice that a variable name can be identical to a statement name.)

In the same way, if you replace the variables _COL_, _ROW_, _TYPE_, and _COEF_ in the previous sparse form data set by COLUMN, ROW, TYPE and COEF, you need the problem definition statements (COL, ROW, TYPE, and COEF) in PROC LP.

```
proc lp sparsedata;
    col column;
    row row;
    type type;
    coef coef;
run;
```

In the sparse form data set, the value $_RHS_$ under the variable $_COL_$ is a special column name, which represents the model's right-hand-side column. If you replace it by a name R, the PROC LP statements would be

```
proc lp sparsedata;
    rhs r;
run;
```

Other special column names for the sparse format are _RHSSEN_ and _RANGE_. The corresponding statements are the RHSSEN and RANGE statements.

PROC LP is case insensitive to variable names and all character values, including the row and column names in the sparse format. The order of the problem definition statements is not important.

For the dense format, a model's row names appear as character values in a SAS data set. For the sparse format, both the row and the column names of the model appear as character values in the data set. Thus, you can put spaces or other special characters in the names. When referring to these names in the problem definition statement or other LP statements, you must use single or double quotes around them. For example, if you replace _RHS_ by R H S in the previous sparse form data set, the PROC LP statements would become

```
proc lp sparsedata;
    rhs "r h s";
run;
```

LP Options

The specifications SPARSEDATA and DATA= in the previous examples are PROC LP options. PROC LP options include

- data set options
- display control options
- interactive control options
- preprocessing options
- branch and bound control options
- sensitivity/parametric/ranging control options
- simplex options

Interactive Processing

Interactive control options include READPAUSE, ENDPAUSE, and so forth. You can run PROC LP interactively using those options. For example, for the blending problem example in the dense form, you can first pause the procedure before iterations start with the READPAUSE option. The PROC LP statements are

proc lp readpause;
run;

When the procedure pauses, you run the PRINT statement to display the initial technological matrix and see if the input is correct. Then you run the PIVOT statement to do one simplex pivot and pause. After that you use the SHOW statement to check the current solution status. Then you apply the RESET statement to tell the procedure to stop as soon as it finds a solution. Now you use the RUN statement to continue the execution. When the procedure stops, you run the PRINT statement again to do a price range analysis and QUIT the procedure. Use SAS %PUT statement to display the contents of PROC LP's macro variable, _ORLP_, which contains iterations and solution information. What follows are the complete statements in batch mode:

```
proc lp readpause;
run;
print matrix(,); /* display all rows and columns. */
pivot;
show status;
reset endpause;
run;
print rangeprice;
quit;
%put &_orlp_;
```

Note: you can force PROC LP to pause during iterations by using the CTRL-BREAK key.

An Integer Programming Example

The following is a simple mixed integer programming problem. Details can be found in Example 3.8 in the "Examples" section.

data;							
input _r	ow_ \$10.	choco	gumdr	ichoco	igumdr	_type_ \$	_rhs_;
dataline	es;						
object	.25	.75	-100	-75	max	•	
cooking	15	40	0	0	le	27000	
color	0	56.25	0	0	le	27000	
package	18.75	0	0	0	le	27000	
condiments	12	50	0	0	le	27000	
chocolate	1	0	-10000	0	le	0	
gum	0	1	0	-10000	le	0	
only_one	0	0	1	1	eq	1	
binary	•	•	1	2	binary	•	
;							

The row with *BINARY* type indicates that this problem is a mixed integer program and all the integer variables are binary. The integer values of the row set an ordering for PROC LP to pick the branching variable when VARSELECT=*PRIOR* is chosen. Smaller values will have higher priorities. The _ROW_ variable here is an alias of the _ID_ variable.

This problem can be solved with the following statements:

```
proc lp canselect=lifo backtrack=obj varselect=far endpause;
run;
quit;
%put &_orlp_;
```

The options CANSELECT=, BACKTRACK=, and VARSELECT= specify the rules for picking the next active problem and the rule to choose the branching variable. In this example, the values *LIFO*, *OBJ* and *FAR* serve as the default values, so the three options can be omitted from the PROC LP statement. The following is the output from the %PUT statement:

```
STATUS=SUCCESSFUL PHASE=3 OBJECTIVE=285 P_FEAS=YES D_FEAS=YES INT_ITER=3
INT_FEAS=2 ACTIVE=0 INT_BEST=285 PHASE1_ITER=1 PHASE2_ITER=5
PHASE3_ITER=5
```

Figure 3.1. The Output of _ORLP_

Preprocessing

Using the PREPROCESS= option, you can apply the preprocessing techniques to pre-solve and then solve the preceding mixed integer program:

```
proc lp preprocess=1 endpause;
run;
quit;
%put &_orlp_;
```

The preprocessing statistics are written to the SAS log file as follows:

```
NOTE: Preprocessing 1 ...
NOTE: 2 upper bounds decreased.
NOTE: 2 coefficients reduced.
NOTE: Preprocessing 2 ...
NOTE: 2 constraints eliminated.
NOTE: Preprocessing done.
```

The new output _ORLP_ is as follows.

```
STATUS=SUCCESSFUL PHASE=3 OBJECTIVE=285 P_FEAS=YES D_FEAS=YES INT_ITER=0
INT_FEAS=1 ACTIVE=0 INT_BEST=285 PHASE1_ITER=1 PHASE2_ITER=4
PHASE3_ITER=0
```

Figure 3.2. The Output of _ORLP_ with Preprocessing Option On

In this example, the number of integer iterations (INT_ITER=) is zero, which means that the preprocessing has reduced the gap between the relaxed linear problem and the mixed integer program to zero.

An MPS Format Conversion Example

If your model input is in MPS input format, you can convert it to the sparse input format of PROC LP using the SAS macro function SASMPSXS. For example, if your have an MPS file called MODEL.MPS and it is stored in the directory C:\OR on a PC, the following program can help you to convert the file and solve the problem.

```
%sasmpsxs(mpsfile="c:\or\model.mps",lpdata=lp);
```

```
data;
   set lp;
   retain i=1;
   if _type_="FREE" and i=1 then
        do;
        _type_="MIN";
        i=0;
        end;
run;
proc lp sparsedata;
run;
```

In the MPS input format, all objective functions, price change rows, and free rows have the type *N*. The SASMPSXS macro marks them as *FREE* rows. After the conversion, you must run a DATA step to identify the objective rows and price change rows. In this example, assume that the problem is one of minimization and the first *FREE* row is an objective row.

Syntax

The following statements are used in PROC LP:

PROC LP options; **COEF** variables; COL variable; ID variable(s); **IPIVOT:** PIVOT; **PRINT** options; QUIT options; RANGE variable ; **RESET** options; RHS variables; **RHSSEN** variables; **ROW** variable(s); RUN: **SHOW** options; **TYPE** variable ; VAR variables;

The TYPE, ID (or ROW), VAR, RHS, RHSSEN, and RANGE statements are used for identifying variables in the problem data set when the model is in the dense input format. In the dense input format, a model's columns appear as variables in the problem data set. The TYPE, ID (or ROW), and RHS statements can be omitted if the input data set contains variables _TYPE_, _ID_ (or _ROW_), and _RHS_; otherwise, they must be used. The VAR statement is optional. When it is omitted, PROC LP treats all numeric variables, which are not explicitly or implicitly included in RHS, RHSSEN, and RANGE statements, as structural variables. The RHSSEN and RANGE statements are optional statements for sensitivity and range analyses. They can be omitted if the input data set contains the _RHSSEN_ and _RANGE_ variables.

The TYPE, COL, ROW (or ID), COEF, RHS, RHSSEN, and RANGE statements are used for identifying variables in the problem data set when the model is in the sparse input format. In the sparse input format, a model's rows and columns appear as observations in the problem data set. The TYPE, COL, ROW (or ID), and COEF statements can be omitted if the input data set contains the _TYPE_ and _COL_ variables, as well as variables beginning with the prefixes _ROW (or _ID) and _COEF. Otherwise, they must be used. The RHS, RHSSEN, and RANGE statements identify the corresponding columns in the model. These statements can be omitted if there are observations that contain the *RHS*, *RHSSEN*, and *RANGE* types or the _RHS_, _RHSSEN_, and _RANGE_ types or the _RHS_, _RHSSEN_, and _RANGE_ names.

The SHOW, RESET, PRINT, QUIT, PIVOT, IPIVOT, and RUN statements are especially useful when executing PROC LP interactively. However, they can also be used in batch mode.

Functional Summary

The options available with PROC LP and it's statements are summarized by purpose in Table 3.1.

Description	Statement	Option
Interactive Statements: perform one integer pivot and pause perform one simplex pivot and pause display solution and related information at cur- rent iteration	IPIVOT PIVOT PRINT	IPIVOT; PIVOT; PRINT <i>options</i> ;
terminate processing immediately reset options specified start or resume optimization show settings of options	QUIT RESET RUN SHOW	QUIT options; RESET options; RUN; SHOW options;
Variable Lists: for the sparse input format, the numeric vari-	COEF	COEF variables;
ables containing the coefficients for the sparse input format, a character vari- able containing the names of a model's	COL	COL variable;
an alias for ROW <i>variable(s)</i> . for the dense input format, the variable con- taining the range constant. For the sparse in- put format, the model's column containing the	ID RANGE	ID variable(s); RANGE variable;
for the dense input format, the variables con- taining the RHS constants. For the sparse in- put format, the model's columns containing	RHS	RHS variables;
for the dense input format, the variables defin- ing RHS change vectors. For the sparse input format, the model's columns representing the	RHSSEN	RHSSEN variables;
for the dense input format, a character variable containing the names of constraints and objec- tive functions. For the sparse input format, the character variables containing the names of a model's rows	ROW	ROW variable(s);
a character variable containing the type of each observation	TYPE	TYPE variables;
for the dense input format, the numeric vari- ables to be treated as structural variables	VAR	VAR variables;

 Table 3.1.
 Functional Summary

Description	Statement	Option
Data Set Options:		
active nodes input data set name	LP	ACTIVEIN=
active nodes output data set name	LP	ACTIVEOUT=
input data set name	LP	DATA=
dual output data set name	LP	DUALOUT=
primal input data set name	LP	PRIMALIN=
primal output data set name	LP	PRIMALOUT=
sparse format data input flag	LP	SPARSEDATA
tableau output data set name	LP	TABLEAUOUT=
Display Control Options:		
display iteration log	LP	FLOW
nonzero tolerance displaying	LP	FUZZ=
inverse of FLOW option	LP	NOFLOW
inverse of PARAPRINT option	LP	NOPARAPRINT
omit some displaying	LP	NOPRINT
inverse of TABLEAUPRINT	LP	NOTABLEAUPRINT
parametric programming displaying	LP	PARAPRINT
inverse of NOPRINT	LP	PRINT
display frequency of iteration or integer	LP	PRINTFREO=
iteration		
level of display desired	LP	PRINTLEVEL =
display the final tableau	LP	TABLEAUPRINT
Interactive Control Options:		
pause before displaying the solution	LP	ENDPAUSE
pause after first feasible solution	LP	FEASIBLEPAUSE
pause frequency of integer solutions	LP	IFEASIBLEPAUSE=
pause frequency of integer iterations	LP	IPAUSE=
inverse of ENDPAUSE	LP	NOENDPAUSE
inverse of FEASIBLEPAUSE	LP	NOFEASIBLEPAUSE
pause frequency of iterations	LP	PAUSE=
pause if within specified proximity	LP	PROXIMITYPAUSE=
pause after data is read	LP	READPAUSE
Preprocessing Options:		
do not perform preprocessing	LP	NOPREPROCESS
perform preprocessing techniques	LP	PREPROCESS
preprocessing error tolerance	LP	PEPSILON=
limit on preprocessing iterations	LP	PMAXIT=
Frefreersmy retailons		
Branch and Bound (BB) Control Options:		
perform automatic node selection technique	LP	AUTO, AUTO(m,n)
backtrack strategy to be used	LP	BACKTRACK=
branch on binary variables first	LP	BINFST

Description	Statement	Option
active node selection strategy	LP	CANSELECT=
a comprehensive node selection control parameter	LP	CONTROL=
a backtrack related technique	LP	DELTAIT=
a measure for pruning BB tree	LP	DOBJECTIVE=
integer tolerance	LP	IEPSILON=
limit on integer iterations	LP	IMAXIT=
a measure for pruning BB tree	LP	IOBJECTIVE=
order of two branched nodes in adding to BB tree	LP	LIFOTYPE=
inverse of AUTO	LP	NOAUTO
inverse of BINFST	LP	NOBINFST
inverse of POSTPROCESS	LP	NOPOSTPROCESS
limits number of branching variables	LP	PENALTYDEPTH=
a measure for pruning BB tree	LP	POBJECTIVE=
perform variables fixing technique	LP	POSTPROCESS
a percentage used in updating WOBJECTIVE	LP	PWOBJECTIVE=
compression algorithm to be used to store ac- tive nodes	LP	TREETYPE=
branching variable selection strategy	LP	VARSELECT=
delays examination of some active nodes	LP	WOBJECTIVE=
Sensitivity/Parametric/Ranging Control Opt	ions:	
inverse of RANGEPRICE	LP	NORANGEPRICE
inverse of RANGERHS	LP	NORANGERHS
limit on perturbing the price vector	LP	PRICEPHI=
range analysis on the price coefficients	LP	RANGEPRICE
range analysis on the RHS vector	LP	RANGERHS
limit on perturbing the RHS vector	LP	RHSPHI=
Simplex Algorithm Control Options:	I D	DEVEN
use devex method		
general error tolerance		EPSILON=
perform goal programming		GUALPROGRAM
largest number used in computation		
reinversion frequency		INVFREQ=
reinversion tolerance		INVIOL=
MAXIT3 and IMAXIT values	LP	MAXII=
limit on phase 1 iterations	LP	MAXIT1=
limit on phase 2 iterations	LP	MAXIT2=
limit on phase 3 iterations	LP	MAXIT3=
inverse of devex	LP	NODEVEX
restore basis after parametric programming	LP	PARARESTORE

Description	Statement	Option
weight of the phase 2 objective function in phase 1	LP	PHASEMIX=
multiple pricing strategy	LP	PRICETYPE=
number of columns to subset in multiple	LP	PRICE=
limit on the number of iterations randomly se- lecting each entering variable during phase 1	LP	RANDOMPRICEMULT=
zero tolerance in ratio test	LP	REPSILON=
scaling type to be performed	LP	SCALE=
zero tolerance in LU decomposition	LP	SMALL=
time pause limit	LP	TIME=
controls the choice of pivot during LU decomposition	LP	U=
RESET Statement Options: The RESET statement supports the same op- tions as the PROC LP statement except for the DATA=, PRIMALIN=, and ACTIVEIN= options, and supports the following additional options.		
new lower bound on listed variables during phase 3	RESET	LOWER (colnames)=
new upper bound on listed variables during phase 3	RESET	UPPER (colnames)=
PRINT statement Options:		
display the best integer solution	PRINT	BEST
display variable summary for the specified columns	PRINT	COLUMN(colnames)
display variable summary and price sensitivity analysis for the specified columns	PRINT	COLUMN(<i>colnames</i>) /SENSITIVITY
display variable summary for the integer variables	PRINT	INTEGER
display variable summary for the nonzero in- teger variables	PRINT	INTEGER_NONZEROS
display variable summary for the integer vari- ables with zero activity	PRINT	INTEGER_ZEROS
display submatrix for the specified rows and columns	PRINT	MATRIX(rownames, colnames)
display formatted submatrix for the specified rows and columns	PRINT	MATRIX(rownames, colnames)/PICTURE
display variable summary for the continuous variables	PRINT	NONINTEGER

Description	Statement	Option
display variable summary for the nonzero con- tinuous variables	PRINT	NONINTEGER_NONZEROS
display variable summary for the variables with nonzero activity	PRINT	NONZEROS
display price sensitivity analysis or price para- metric programming	PRINT	PRICESEN
display price range analysis	PRINT	RANGEPRICE
display RHS range analysis	PRINT	RANGERHS
display RHS sensitivity analysis or RHS para- metric programming	PRINT	RHSSEN
display constraint summary for the specified rows	PRINT	ROW(rownames)
display constraint summary and RHS sensitivity	PRINT	ROW(rownames)
analysis for the specified rows		/SENSITIVITY
display solution, variable, and constraint summaries	PRINT	SOLUTION
display current tableau	PRINT	TABLEAU
display variables with zero activity	PRINT	ZEROS
SHOW Statement Options:		
display options applied	SHOW	OPTIONS
display status of the current solution	SHOW	STATUS
QUIT Statement Option:		
terminate PROC LP	QUIT	QUIT
save the defined output data sets and then terminate PROC LP	QUIT	QUIT/SAVE

PROC LP Statement

PROC LP options;

This statement invokes the procedure. The following options can appear in the PROC LP statement.

Data Set Options

ACTIVEIN=SAS-data-set

names the SAS data set containing the active nodes in a branch and bound tree that is to be used to restart an integer program.

ACTIVEOUT=SAS-data-set

names the SAS data set in which to save the current branch and bound tree of active nodes.

DATA=SAS-data-set

names the SAS data set containing the problem data. If the DATA= option is not specified, PROC LP uses the most recently created SAS data set.

DUALOUT=SAS-data-set

names the SAS data set that contains the current dual solution (shadow prices) on termination of PROC LP. This data set contains the current dual solution only if PROC LP terminates successfully.

PRIMALIN=SAS-data-set

names the SAS data set that contains a feasible solution to the problem defined by the DATA= data set. The data set specified in the PRIMALIN= option should have the same format as a data set saved using the PRIMALOUT= option. Specifying the PRIMALIN= option is particularly useful for continuing iteration on a problem previously attempted. It is also useful for performing sensitivity analysis on a previously solved problem.

PRIMALOUT=SAS-data-set

names the SAS data set that contains the current primal solution when PROC LP terminates.

SPARSEDATA

tells PROC LP that the data are in the sparse input format. If this option is not specified, PROC LP assumes that the data are in the dense input format. See the "Sparse Data Input Format" section on page 91 for information about the sparse input format.

TABLEAUOUT=SAS-data-set

names the SAS data set in which to save the final tableau.

Display Control Options

FLOW

requests that a journal (the Iteration Log) of pivot information be displayed at each PRINTFREQ= iteration. This includes the names of the variables entering and leaving the basis, the reduced cost of the entering variable, and the current objective value.

FUZZ=e

displays all numbers within e of zero as zeros. The default value is 1E - 10.

NOFLOW

is the inverse of the FLOW option.

NOPARAPRINT

is the inverse of the PARAPRINT option.

NOPRINT

suppresses the display of the Variable, Constraint, and Sensitivity Analysis summaries. This option is equivalent to the PRINTLEVEL=0 option.

NOTABLEAUPRINT

is the inverse of the TABLEAUPRINT option.

PARAPRINT

indicates that the solution be displayed at each pivot when performing parametric programming.

PRINT

is the inverse of the NOPRINT option.

PRINTFREQ=m

indicates that after every mth iteration, a line in the (Integer) Iteration Log be displayed. The default value is 1.

PRINTLEVEL=*i*

indicates the amount of displaying that the procedure should perform.

PRINTLEVEL=-2	only messages to the SAS log are displayed.
PRINTLEVEL=-1	is equivalent to NOPRINT unless the problem is infeasi- ble. If it is infeasible, the infeasible rows are displayed in the Constraint Summary along with the Infeasible Infor- mation Summary.
PRINTLEVEL=0	is identical to NOPRINT.
PRINTLEVEL=1	all output is displayed.

The default value is 1.

TABLEAUPRINT

indicates that the final tableau be displayed.

Interactive Control Options

ENDPAUSE

requests that PROC LP pause before displaying the solution. When this pause occurs, you can enter the RESET, SHOW, PRINT, RUN, and QUIT statements.

FEASIBLEPAUSE

requests that PROC LP pause after a feasible (not necessarily integer feasible) solution has been found. At a pause, you can enter RESET, SHOW, PRINT, PIVOT, RUN, and QUIT statements.

IFEASIBLEPAUSE=n

requests that PROC LP pause after every IFEASIBLEPAUSE= integer feasible solutions. At a pause, you can enter RESET, SHOW, PRINT, IPIVOT, PIVOT, RUN, and QUIT statements. The default value is 99999999.

IPAUSE=n

requests that PROC LP pause after every *n* integer iterations. At a pause, you can enter the RESET, SHOW, PRINT, IPIVOT, PIVOT, RUN, and QUIT statements. The default value is 999999999.

NOENDPAUSE

is the inverse of the ENDPAUSE option.

NOFEASIBLEPAUSE

is the inverse of the FEASIBLEPAUSE option.

PAUSE=n

requests that PROC LP pause after every *n* iterations. At a pause, you can enter the RESET, SHOW, PRINT, IPIVOT, PIVOT, RUN, and QUIT statements. The default value is 999999999.

PROXIMITYPAUSE=r

This option causes the procedure to pause if at least one integer feasible solution has been found and the objective value of the current best integer solution can be determined to be within *r* units of the optimal integer solution. This distance, called proximity, is also displayed on the Integer Iteration Log. Note that the proximity is calculated using the minimum (maximum if the problem is maximization) objective value among the nodes that remain to be explored in the branch and bound tree as a bound on the value of the optimal integer solution. Following the first PROX-IMITYPAUSE= pause, in order to avoid a pause at every iteration thereafter, it is recommended that you reduce this measure through the use of a RESET statement. Otherwise, if any other option or statement that causes the procedure to pause is used while the PROXIMITYPAUSE= option is in effect, pause interferences may occur. When this pause occurs, you can enter RESET, SHOW, PRINT, IPIVOT, PIVOT, RUN, and QUIT statements. The default value is 0.

READPAUSE

requests that PROC LP pause after the data have been read and the initial basis inverted. When this pause occurs, you can enter RESET, SHOW, PRINT, IPIVOT, PIVOT, RUN, or QUIT statements.

Preprocessing Control Options

NOPREPROCESS

is the inverse of the PREPROCESS option.

PREPROCESS

perform preprocessing techniques. See the "Preprocessing" section on page 99 for further discussion.

PEPSILON=*e*

specifies a positive number close to zero. This value is an error tolerance in the preprocessing. If the value is too small, any marginal changes may cause the preprocessing to repeat itself. However, if the value is too large, it may alter the optimal solution or falsely claim that the problem is infeasible. The default value is 1E - 8.

PMAXIT=n

performs at most *n* preprocessings. Preprocessing repeats itself if it improves some bounds or fixes some variables. However when a problem is large and dense, each preprocessing may take a significant amount of CPU time. This option limits the number of preprocessings PROC LP performs. It can also reduce the build-up of round-off errors. The default value is 100.

Branch and Bound Algorithm Control Options

AUTO, AUTO(m,n)

This option automatically sets and adjusts the value of the CONTROL= option. Initially, it sets CONTROL=0.70 concentrating on finding an integer feasible solution or an upper bound. When an upper bound is found, it sets CONTROL=0.5 concentrating on efficiency and lower bound improvement. When the number of active problems exceeds m, it starts to increase the value of CONTROL= gradually to keep the size of active problems under control. When total active problems exceed n, CONTROL=1 will keep the active problems from further growing. You can alter the automatic process by resetting the value of CONTROL= interactively.

The default values of m and n are 20000 and 250000, respectively. You can change the two values according to your computer's space and memory capacities.

BACKTRACK=rule

specifies the rule used to choose the next active problem when backtracking is required. One of the following can be specified:

- BACKTRACK=LIFO
- BACKTRACK=FIFO
- BACKTRACK=OBJ
- BACKTRACK=PROJECT
- BACKTRACK=PSEUDOC
- BACKTRACK=ERROR

The default value is *OBJ*. See the "Integer Programming" section on page 99 for further discussion.

BINFST

requests that PROC LP branch on binary variables first when integer and binary variables present. The reasoning behind this is that a subproblem will usually be fathomed or found integer feasible after less than 20% of its variables have been fixed. Considering binary variables first attempts to reduce the size of the branch and bound tree. It is a heuristic technique.

CANSELECT=*rule*

specifies the rule used to choose the next active problem when backtracking is not required or used. One of the following can be specified:

- CANSELECT=LIFO
- CANSELECT=FIFO
- CANSELECT=OBJ
- CANSELECT=PROJECT
- CANSELECT=PSEUDOC
- CANSELECT=ERROR

The default value is *LIFO*. See the "Integer Programming" section on page 99 for further discussion.

CONTROL=r

specifies a number between 0 and 1. This option combines CANSELECT= and other rules to choose the next active problem. It takes into consideration of three factors, efficiency, improving lower and upper bounds. When r is close to 0, PROC LP concentrates on improving lower bound (upper bound for maximization). However, the efficiency per integer iteration is usually the worst. When r is close to 1, PROC LP concentrates on improving upper bound (lower bound for maximization). In addition, the growth of active problems will be controlled and stopped at r = 1. When its value is around 0.5, PROC LP will be in the most efficient state in terms of CPU time and integer number of iterations. CONTROL= option will be automatically adjusted when the AUTO option is applied.

DELTAIT=r

This option is used to modify the exploration of the branch and bound tree. If more than r integer iterations have occurred since the last integer solution was found, then the procedure uses the backtrack strategy in choosing the next node to be explored. The default value is 3 times the number of integer variables.

DOBJECTIVE=r

specifies that PROC LP should discard active nodes that cannot lead to an integer solution with the objective at least as small (or as large for maximizations) as the objective of the relaxed problem plus (minus) *r*. The default value is $+\infty$.

IEPSILON=e

requests that PROC LP consider an integer variable as having an integer value if its value is within e units of an integer. The default value is 1E - 7.

IMAXIT=n

performs at most *n* integer iterations. The default value is 100.

IOBJECTIVE=r

specifies that PROC LP should discard active nodes unless the node could lead to an integer solution with the objective smaller (or larger for maximizations) than *r*. The default value is $+\infty$ for minimization ($-\infty$ for maximization).

LIFOTYPE=*i*

specifies the order in which to add the two of the newly branched active nodes to the LIFO list.

LIFOTYPE=0	add the node with minimum penalty first
LIFOTYPE=1	add the node with maximum penalty first
LIFOTYPE=2	add the node resulting from adding $x_i \geq \lceil x^{opt}(k)_i \rceil$ first
LIFOTYPE=3	add the node resulting from adding $x_i \leq \lfloor x^{opt}(k)_i \rfloor$ first

The default value is 0.

NOAUTO

turn off AUTO option.

NOBINFST

turn off BINFST option.

NOPOSTPROCESS

do not perform postprocesing.

PENALTYDEPTH=*m*

requests that PROC LP examine *m* variables as branching candidates when VARS-ELECT=*PENALTY*. If the PENALTYDEPTH= option is not specified when VARS-ELECT=*PENALTY*, then all of the variables are considered branching candidates. The default value is the number of integer variables. See the "Integer Programming" section on page 99 for further discussion.

POBJECTIVE=*r*

specifies that PROC LP should discard active nodes that cannot lead to an integer solution with objective at least as small as $o + |o| \times \text{POBJECTIVE}=$ (at least as large as $o - |o| \times \text{POBJECTIVE}=$ for maximizations) where *o* is the objective of the relaxed noninteger constrained problem. The default value is $+\infty$.

POSTPROCESS

attempts to fix binary variables globally based on the relationships among the reduced cost and objective value of the relaxed problem and the objective value of current best integer feasible solution.

PWOBJECTIVE=*r*

This option gives a percentage for use in the automatic update of the WOBJECTIVE= option. If the WOBJECTIVE= option is not specified in PROC LP, then when an integer feasible solution is found, the value of the option is updated to be $b + q \times r$ where b is the best bound on the value of the optimal integer solution and q is the current proximity. Note that for maximizations, $b - q \times r$ is used. The default value is 0.95.

TREETYPE=*i*

specifies a data compression algorithm.

TREETYPE=0	no data compression
TREETYPE=1	Huffman coding compression routines
TREETYPE=2	adaptive Huffman coding compression routines
TREETYPE=3	adaptive arithmetic coding compression routines

For IP or MIP problems, the basis and bounds information of each active node is saved to a utility file. When the number of active nodes increases, the size of the utility file becomes larger and larger. If PROC LP runs into a disk problem, like "disk full ..." or "writing failure ...", you can use this option to compress the utility file. For more information on the data compression routines, refer to Nelton (1992). The default value is 0.

VARSELECT=rule

specifies the rule used to choose the branching variable on an integer iteration.

- VARSELECT=CLOSE
- VARSELECT=PRIOR
- VARSELECT=PSEUDOC
- VARSELECT=FAR
- VARSELECT=PRICE
- VARSELECT=*PENALTY*

The default value is *FAR*. See the "Integer Programming" section on page 99 for further discussion.

WOBJECTIVE=r

specifies that PROC LP should delay examination of active nodes that cannot lead to an integer solution with objective at least as small (as large for maximizations) as *r*, until all other active nodes have been explored. The default value is $+\infty$ for minimization ($-\infty$ for maximization).

Sensitivity/Parametric/Ranging Control Options

NORANGEPRICE

is the inverse of the RANGEPRICE option.

NORANGERHS

is the inverse of the RANGERHS option.

$\textbf{PRICEPHI}{=}\Phi$

specifies the limit for parametric programming when perturbing the price vector. See the "Parametric Programming" section on page 110 for further discussion. See Example 3.5 for an illustration of this option.

RANGEPRICE

indicates that range analysis is to be performed on the price coefficients. See the "Range Analysis" section on page 110 for further discussion.

RANGERHS

indicates that range analysis is to be performed on the right-hand-side vector. See the "Range Analysis" section on page 110 for further discussion.

$RHSPHI=\Phi$

specifies the limit for parametric programming when perturbing the right-hand-side vector. See the "Parametric Programming" section on page 110 for further discussion.

Simplex Algorithm Control Options

DEVEX

indicates that the devex method of weighting the reduced costs be used in pricing (Harris 1975).

EPSILON=e

specifies a positive number close to zero. It is used in the following instances:

During phase 1, if the sum of the basic artificial variables is within e of zero, the current solution is considered feasible. If this sum is not exactly zero, then there are artificial variables within e of zero in the current solution. In this case, a note is displayed on the SAS log.

During phase 1, if all reduced costs are $\leq e$ for nonbasic variables at their lower bounds and $\geq e$ for nonbasic variables at their upper bounds and the sum of infeasibilities is greater than e, then the problem is considered infeasible. If the maximum reduced cost is within e of zero, a note is displayed on the SAS log.

During phase 2, if all reduced costs are $\leq e$ for nonbasic variables at their lower bounds and $\geq e$ for nonbasic variables at their upper bounds, then the current solution is considered optimal.

During phases 1, 2, and 3, the EPSILON= option is also used to test if the denominator is different from zero before performing the ratio test to determine which basic variable should leave the basis.

The default value 1E - 8.

GOALPROGRAM

specifies that multiple objectives in the input data set are to be treated as sequential objectives in a goal-programming model. The value of the right-hand-side variable in the objective row gives the priority of the objective. Lower numbers have higher priority.

INFINITY=r

specifies the largest number PROC LP uses in computation. The INFINITY= option is used to determine when a problem has an unbounded variable value. The default value is the largest double precision number. *

INVFREQ=m

reinverts the current basis matrix after m major and minor iterations. The default value is 100.

INVTOL=r

reinverts the current basis matrix if the largest element in absolute value in the decomposed basis matrix is greater than r. If after reinversion this condition still holds, then the value of the INVTOL= option is increased by a factor of 10 and a note indicating this modification is displayed on the SAS log. When r is frequently exceeded, this may be an indication of a numerically unstable problem. The default value is 1000.

*This value is system dependent.

MAXIT=n

simultaneously sets the values of the MAXIT1=, MAXIT2=, MAXIT3=, and IMAXIT= options.

MAXIT1=n

performs at most $n \ge 0$ phase 1 iterations. The default value is 100.

MAXIT2=n

performs at most $n \ge 0$ phase 2 iterations. If MAXIT2=0, then only phase 1 is entered so that on successful termination PROC LP will have found a feasible, but not necessarily optimal, solution. The default value is 100.

MAXIT3=n

performs at most $n \ge 0$ phase 3 iterations. All dual pivots are counted as phase 3 pivots. The default value is 99999999.

NODEVEX

is the inverse of the DEVEX option.

PARARESTORE

indicates that following a parametric programming analysis, PROC LP should restore the basis.

PHASEMIX=r

specifies a number between 0 and 1. When the number is positive, PROC LP tries to improve the objective function of phase 2 during phase 1. The PHASEMIX= option is a weight factor of the phase 2 objective function in phase 1. The default value is 0.

PRICETYPE=pricetype

specifies the type of multiple pricing to be performed. If this option is specified and the PRICE= option is not specified, then PRICE= is assumed to be 10.

- PRICETYPE=*COMPLETE*
- PRICETYPE=DYNAMIC
- PRICETYPE=NONE
- PRICETYPE=PARTIAL

The default value is *PARTIAL*. See the "Pricing" section on page 97 for a description of this process.

PRICE=m

specifies the number of columns to subset when multiple pricing is used in selecting the column to enter the basis (Greenberg 1978). The type of suboptimization used is determined by the PRICETYPE= option. See the "Pricing" section on page 97 for a description of this process.

RANDOMPRICEMULT=r

specifies a number between 0 and 1. This option sets a limit, in phase 1, on the number of iterations when PROC LP will randomly pick the entering variables. The limit equals r times the number of nonbasic variables, or the number of basic variables, which ever is smaller. The default value of the RANDOMPRICEMULT= option is 0.01.

REPSILON=*e*

specifies a positive number close to zero. The REPSILON= option is used in the ratio test to determine which basic variable is to leave the basis. The default value is 1E - 10.

SCALE=scale

specifies the type of scaling to be used.

- SCALE=NONE
- SCALE=COLUMN
- SCALE=ROW
- SCALE=BOTH

The default value is *NONE*. See the "Scaling" section on page 98 for further discussion.

SMALL=e

specifies a positive number close to zero. Any element in a matrix with a value less than e is set to zero. The default value is machine dependent.

TIME=t

checks at each iteration to see if t seconds have elapsed since PROC LP began. If more than t seconds have elapsed, the procedure pauses and displayed the current solution. The default value is 120 seconds.

U=r

enables PROC LP to control the choice of pivots during LU decomposition and updating the basis matrix. The variable r should take values between EPSILON and 1.0 because small values of r bias the algorithm toward maintaining sparsity at the expense of numerical stability and vice versa. The more sparse the decomposed basis is, the less time each iteration takes. The default value is 0.1.

COEF Statement

COEF variables;

For the sparse input format, the COEF statement specifies the numeric variables in the problem data set that contain the coefficients in the model. The value of the coefficient variable in a given observation is the value of the coefficient in the column and row specified in the COLUMN and ROW variables in that observation. For multiple ROW variables, the LP procedure maps the ROW variables to the COEF variables on the basis of their order in the COEF and ROW statements. There must be the same number of COEF variables as ROW variables. If the COEF statement is omitted, the procedure looks for the default variable names that have the prefix _COEF.

COL Statement

COL variable;

For the sparse input format, the COL statement specifies a character variable in the problem data set that contains the names of the columns in the model. Columns in the model are either structural variables, right-hand-side vectors, right-hand-side change vectors, or a range vector. The COL variable must be a character variable. If the COL statement is omitted, the LP procedure looks for the default variable name _COL_.

ID Statement

ID variable(s);

For the dense input format, the ID statement specifies a character variable in the problem data set that contains a name for each constraint coefficients row, objective coefficients row, and variable definition row. If the ID statement is omitted, the LP procedure looks for the default variable name, $_ID_$. If this variable is not in the problem data set, the procedure assigns the default name $_OBSxx_$ to each row, where xx specifies the observation number in the problem data set.

For the sparse input format, the ID statement specifies the character variables in the problem data set that contain the names of the rows in the model. Rows in the model are one of the following types: constraints, objective functions, bounding rows, or variable describing rows. The ID variables must be character variables. There must be the same number of ID variables as variables specified in the COEF statement. If the ID statement is omitted, the LP procedure looks for the default variable names having the prefix _ID.

Note: The ID statement is an alias for the ROW statement.

IPIVOT Statement

IPIVOT;

The IPIVOT statement causes the LP procedure to execute one integer branch and bound pivot and pause. If you use the IPIVOT statement while the PROXIMITY-PAUSE option is in effect, pause interferences may occur. To avoid such interferences, you must either reset the PROXIMITYPAUSE value or submit IPIVOT; RUN; instead of IPIVOT;.

PIVOT Statement

PIVOT;

The PIVOT statement causes the LP procedure to execute one simplex pivot and pause.

PRINT Statement

PRINT options;

The PRINT statement is useful for displaying part of a solution summary, examining intermediate tableaus, performing sensitivity analysis, and using parametric programming. In the options, the **colnames** and **rownames** lists can be empty, in which case the LP procedure displays tables with all columns or rows, or both. If a column or a row name has spaces or other special characters in it, the name must be enclosed in single or double quotes when it appears in the argument. The options that can be used with this statement are as follows.

BEST

displays a Solution, Variable, and Constraint Summary for the best integer solution found.

COLUMN (colnames)/SENSITIVITY

displays a Variable Summary containing the logical and structural variables listed in the colnames list. If the /SENSITIVITY option is included, then sensitivity analysis is performed on the price coefficients for the listed colnames structural variables.

INTEGER

displays a Variable Summary containing only the integer variables.

INTEGER_NONZEROS

displays a Variable Summary containing only the integer variables with nonzero activity.

INTEGER_ZEROS

displays a Variable Summary containing only the integer variables with zero activity.

MATRIX (*rownames*,*colnames*)/**PICTURE**

displays the submatrix of the matrix of constraint coefficients defined by the rownames and colnames lists. If the /PICTURE option is included, then the formatted submatrix is displayed. The format used is summarized in Table 3.2.

	Condi	ition	on the Co	peffic	eient x		Symbols	Printed
			abs(x)	=	0			
0		<	abs(x)	<		.000001	sgn(x)	"Z"
	.000001	\leq	abs(x)	<		.00001	sgn(x)	"Y"
	.00001	\leq	abs(x)	<		.0001	sgn(x)	"X"
	.0001	\leq	abs(x)	<		.001	sgn(x)	"W"
	.001	\leq	abs(x)	<		.01	sgn(x)	"V"
	.01	\leq	abs(x)	<		.1	sgn(x)	"U"
	.1	\leq	abs(x)	<	1		sgn(x)	"T"
			abs(x)	=	1		sgn(x)	"1"
1		<	abs(x)	<	10		sgn(x)	"A"
10		\leq	abs(x)	<	100		sgn(x)	"B"
100		\leq	abs(x)	<	1000		sgn(x)	"C"
1000		\leq	abs(x)	<	10000		sgn(x)	"D"
10000		\leq	abs(x)	<	100000		sgn(x)	"E"
100000		\leq	abs(x)	<	1.0E06		sgn(x)	"F"

Table 3.2. Format Summary

NONINTEGER

displays a variable summary containing only the continuous variables.

NONINTEGER_NONZEROS

displays a variable summary containing only the continuous variables with nonzero activity.

NONZEROS

displays a variable summary containing only the variables with nonzero activity.

PRICESEN

displays the results of parametric programming for the current value of the PRICEPHI option, the price coefficients, and all of the price change vectors.

RANGEPRICE

performs range analysis on the price coefficients.

RANGERHS

performs range analysis on the right-hand-side vector.

RHSSEN

displays the results of parametric programming for the current value of the RHSPHI option, the right-hand-side coefficients, and all of the right-hand-side change vectors.

ROW (*rownames***)/SENSITIVITY**

displays a constraint summary containing the rows listed in the rowname list. If the /SENSITIVITY option is included, then sensitivity analysis is performed on the right-hand-side coefficients for the listed rownames.

SOLUTION

displays the solution summary, including the variable summary and the constraint summary.

TABLEAU

displays the current tableau.

ZEROS

displays a variable summary containing only the variables with zero activity. This may be useful in the analysis of ON/OFF, ZERO/ONE, scheduling, and assignment applications.

QUIT Statement

QUIT options;

The QUIT statement causes the LP procedure to terminate processing immediately. No further displaying is performed and no output data sets are created.

The QUIT/SAVE statement causes the LP procedure to save the output data sets, defined in the PROC LP statement or in the RESET statement, and then terminate the procedure.

RANGE Statement

RANGE variable;

For the dense input format, the RANGE statement identifies the variable in the problem data set that contains the range coefficients. These coefficients enable you to specify the feasible range of a row. For example, if the *i*th row is

$$a^T x \leq b_i$$

and the range coefficient for this row is $r_i > 0$, then all values of x that satisfy

$$b_i - r_i \le a^T x \le b_i$$

are feasible for this row. Table 3.3 shows the bounds on a row as a function of the row type and the sign on a nonmissing range coefficient r.

		Bounds				
r	_TYPE_	Lower	Upper			
$\neq 0$	LE	$b - \mid r \mid$	b			
$\neq 0$	GE	b	$b+\mid r\mid$			
> 0	EQ	b	b+r			
< 0	EQ	b+r	b			

Table 3.3. Interpretation of the Range Coefficient

If you include a range variable in the model and have a missing value or zero for it in a constraint row, then that constraint is treated as if no range variable had been included.

If the RANGE statement is omitted, the LP procedure assumes that the variable named _RANGE_ contains the range coefficients.

For the sparse input format, the RANGE statement gives the name of a column in the problem data set that contains the range constants. If the RANGE statement is omitted, then the LP procedure assumes that the column named _RANGE_ or the column with the *RANGE* keyword in the problem data set contains the range constants.

RESET Statement

RESET options;

The RESET statement is used to change options after the LP procedure has started execution. All of the options that can be set in the PROC LP statement can also be reset with the RESET statement, except for the DATA=, the PRIMALIN=, and the ACTIVEIN= options. In addition to the options available with the PROC LP statement, the following two options can be used.

LOWER(colnames)=n;

during phase 3, this sets the lower bound on all of the structural variables listed in the colnames list to an integer value *n*. This may contaminate the branch and bound tree. All nodes that descend from the current problem have lower bounds that may be different from those input in the problem data set.

UPPER(*colnames*)=*n*;

during phase 3, this sets the upper bound on all of the structural variables listed in the colnames list to an integer value *n*. This may contaminate the branch and bound tree. All nodes that descend from the current problem have upper bounds that may be different from those input in the problem data set.

RHS Statement

RHS variables;

For the dense input format, the RHS statement identifies variables in the problem data set that contain the right-hand-side constants of the linear program. Only numeric variables can be specified. If more than one variable is included in the RHS statement, the LP procedure assumes that problems for several linear programs are defined in the problem data set. A new linear program is defined for each variable in the RHS list. If the RHS statement is omitted, the procedure assumes that the variable named _RHS_ contains the right-hand-side constants.

For the sparse input format, the RHS statement gives the names of one or more columns in the problem data set that are to be considered as right-hand-side constants. If the RHS statement is omitted, then the LP procedure assumes that the column named _RHS_ or columns with the *RHS* keyword in the problem data set contain the right-hand-side constants. See the "Sparse Data Input Format" section on page 91 for further discussion.

In the default, the LP procedure assumes that the RHS constant is a zero vector for the dense and sparse input formats.

RHSSEN Statement

RHSSEN variables;

For the dense input format, the RHSSEN statement identifies variables in the problem data set that define change vectors for examining the sensitivity of the optimal solution to changes in the RHS constants. If the RHSSEN statement is omitted, then the LP procedure assumes that the variable named _RHSSEN_ contains a right-hand-side change vector.

For the sparse input format, the RHSSEN statement gives the names of one or more columns in the problem data set that are to be considered as change vectors. If the RHSSEN statement is omitted, then the LP procedure assumes that the column named __RHSSEN_ or columns with the *RHSSEN* keyword in the problem data set contain the right-hand-side change vectors. For further information, see the "Sparse Data Input Format" section on page 91, the "Right-Hand-Side Sensitivity Analysis" section on page 108, and the "Right-Hand-Side Parametric Programming" section on page 111.

ROW Statement

ROW variable(s);

For the dense input format, the ROW statement specifies a character variable in the problem data set that contains a name for each row of constraint coefficients, each row of objective coefficients and each variable describing row. If the ROW statement is omitted, the LP procedure looks for the default variable name, $_ROW_$. If this is not such a variable in the problem data set, the procedure assigns the default name $_OBSxx_$ to each row, where xx specifies the observation number in the problem data set.

For the sparse input format, the ROW statement specifies the character variables in the problem data set that contain the names of the rows in the model. Rows in the model are one of the following types: constraints, objective functions, bounding rows, or variable describing rows. The ROW variables must be character variables. There must be the same number of ROW variables as variables specified in the COEF statement. If the ROW statement is omitted, the LP procedure looks for the default variable names having the prefix _ROW.

RUN Statement

RUN;

The RUN statement causes optimization to be started or resumed. The TITLE or OPTIONS statement should not appear between PROC LP and RUN statements.

SHOW Statement

SHOW options;

The SHOW statement specifies that the LP procedure display either the *current options* or the *current solution status* on the SAS log.

OPTIONS

requests that the current options be displayed on the SAS log.

STATUS

requests that the status of the current solution be displayed on the SAS log.

TYPE Statement

TYPE variable;

The TYPE statement specifies a character variable in the problem data set that contains the type identifier for each observation. This variable has keyword values that specify how the LP procedure should interpret the observation. If the TYPE statement is omitted, the procedure assumes that the variable named _TYPE_ contains the type keywords.

For the dense input format, the type variable identifies the constraint and objective rows and rows that contain information about the variables. The type variable should have nonmissing values in all observations.

For the sparse input format, the type variable identifies a model's rows and columns. In an observation, a nonmissing type is associated with either a row or a column. If there are many columns sharing the same type, you can define a row of that type. Then, any nonmissing values in that row set the types of the corresponding columns.

The following are valid values for the TYPE variable in an observation:

MIN	contains the price coefficients of an objective row, for example, c in the problem (mip) , to be minimized.
MAX	contains the price coefficients of an objective row, for example, c , to be maximized.
EQ (=)	contains coefficients of an equality constrained row.
LE (≤)	contains coefficients of an inequality, less than or equal to, constrained row.
GE (≥)	contains coefficients of an inequality, greater than or equal to, con- strained row.
SOSEQ	identifies the row as specifying a special ordered set. The variables flagged in this row are members of a set exactly one of which must be above its lower bound in the optimal solution. Note that variables in this type of special ordered set must be integer .
SOSLE	identifies the row as specifying a special ordered set. The variables flagged in this row are members of a set in which only one can be above its lower bound in the optimal solution.
UNRSTRCT	identifies those structural variables to be considered as unre- stricted variables. These are variables for which ℓ_i =- <i>INFINITY</i> and u_i =+ <i>INFINITY</i> . Any variable that has a 1 in this observation is con- sidered an unrestricted variable.
LOWERBD	identifies lower bounds on the structural variables. If all structural
---------	--
	variables are to be nonnegative, that is, $\ell_i = 0$, then you do not
	need to include an observation with the LOWERBD keyword in
	a variable specified in the TYPE statement. Missing values for
	variables in a lower-bound row indicate that the variable has zero
	lower bound.

Note: A variable with lower or upper bounds cannot be identified as unrestricted.

- UPPERBD identifies upper bounds u_i on the structural variables. For each structural variable that is to have an upper bound $u_i = +INFINITY$, the observation must contain a missing value or the current value of *INFINITY*. All other values are interpreted as upper bounds, including 0.
- *FIXED* identifies variables that have fixed values. A nonmissing value in a row with *FIXED* type keyword gives the constant value of that variable.
- *INTEGER* identifies variables that are integer constrained. In a feasible solution, these variables must have integer values. A missing value in a row with *INTEGER* type keyword indicates that that variable is not integer constrained. The value of variables in the *INTEGER* row gives an ordering to the integer-constrained variables that is used when the VARSELECT= option equals *PRIOR*.

Note: Every integer-constrained variable must have an upper bound defined in a row with type *UPPERBD*. See the "Controlling the Branch and Bound Search" section on page 102 for further discussion.

- *BINARY* identifies variables that are constrained to be either 0 or 1. This is equivalent to specifying that the variable is an integer variable and has a lower bound of 0 and an upper bound of 1. A missing value in a row with *BINARY* type keyword indicates that that variable is not constrained to be 0 or 1. The value of variables in the *BINARY* row gives an ordering to the integer-constrained variables that is used when the VARSELECT= option equals *PRIOR*. See the "Controlling the Branch and Bound Search" section on page 102 for further discussion.
- BASIC identifies variables that form an initial basic feasible solution. A missing value in a row with BASIC type indicates that that variable is not basic.
- PRICESENidentifies a vector that is used to evaluate the sensitivity of the op-
timal solution to changes in the objective function. See the "Price
Sensitivity Analysis" section on page 109 and the "Price Paramet-
ric Programming" section on page 111 for further discussion.
- *FREE* identifies a nonbinding constraint. Any number of *FREE* constraints can appear in a problem data set.

RHS	identifies a right-hand-side column in the sparse input format. This will replace the RHS statement. It is useful when converting the MPS format into the sparse format of PROC LP. See the "Converting MPS Format" section on page 93 for more information.
RHSSEN	identifies a right-hand-side sensitivity analysis vector in the sparse input format. This replaces the RHSSEN statement. It is useful when converting the MPS format into the sparse format of PROC LP. See the "Converting MPS Format" section on page 93 for more information.
RANGE	identifies a range vector in the sparse input format. This replaces the RANGE statement. It is useful when converting the MPS for- mat into the sparse format of PROC LP. See the "Converting MPS Format" section on page 93 for more information.

VAR Statement

VAR variables;

For the dense input format, the VAR statement identifies variables in the problem data set that are to be interpreted as structural variables in the linear program. Only numeric variables can be specified. If no VAR statement is specified, the LP procedure uses all numeric variables not included in an RHS or RHSSEN statement as structural variables.

Details

Missing Values

The LP procedure treats missing values as missing in all rows except those that identify either upper or lower bounds on structural variables. If the row is an upper-bound row, then the type identifier is *UPPERBD* and the LP procedure treats missing values as *+INFINITY*. If the row is a lower-bound row, then the type identifier is *LOWERBD* and the LP procedure treats the missing values as 0, except for the variables that are identified as *UNRSTRCT*.

Sparse Data Input Format

The sparse format to PROC LP is designed to enable you to specify only the nonzero coefficients in the description of linear programs, integer programs, and mixed-integer programs. The SAS data set that describes the sparse model must contain at least four SAS variables:

- a type variable
- a column variable
- a row variable
- a coefficient variable

Each observation in the data set associates a type with a row or a column, and defines a coefficient or a numerical value in the model. In addition to the keywords in the dense format, PROC LP also recognizes the keywords *RHS*, *RHSSEN*, and *RANGE* as values of the type variable. The values of the row and column variables are the names of the rows and columns in the model. The values of the coefficient variables specify the coefficients or other numerical data. The SAS data set can contain multiple pairs of row and coefficient variables. In this way, more information about the model can be specified in each observation in the data set. See Example 3.2 for details.

Table 3.4 shows the keywords that are recognized by PROC LP and in which variables can appear in the problem data set. The SAS data set that describes the sparse model must contain at least four SAS variables: a type variable, a column variable, a row variable, and a coefficient variable. Each observation in the data set defines one or more rows or coefficients in the model. The value of the type variable is a keyword that tells PROC LP how to interpret the observation. The values of the row and column variables name the rows and columns in the model. The values of the coefficient variables define coefficients and lower and upper bounds and identify model variables with type BASIC, FIXED, BINARY, and INTEGER. All character values in the sparse data input format are case insensitive.

TYPE (_TYPE_)	COL (_COL_)
MIN	
MAX	
EQ	
LE	
GE	
SOSEQ	
SOSLE	
UNRSTRCT	
LOWERBD	
UPPERBD	
FIXED	
INTEGER	
BINARY	
BASIC	
PRICESEN	
FREE	
RHS	_RHS_
RHSSEN	_RHSSEN_
RANGE	_RANGE_
*XXXXXXX	

Table 3.4. Variable Keywords Used in the Problem Data Set

Follow these rules for sparse data input:

- The order of the observations is unimportant.
- Each unique column name appearing in the COL variable defines a unique column in the model.
- Each unique row name appearing in the ROW variable defines a unique row in the model.
- The type of the row is identified when an observation in which the row name appears (in a ROW variable) has type *MIN*, *MAX*, *LE*, *GE*, *EQ*, *SOSLE*, *SOSEQ*, *LOWERBD*, *UPPERBD*, *UNRSTRCT*, *FIXED*, *BINARY*, *INTEGER*, *BASIC*, *FREE*, or *PRICESEN*.
- The type of each row must be identified at least once. If a row is given a type more than once, the multiple definitions must be identical.
- When there are multiple rows named in an observation (that is, when there are multiple ROW variables), the TYPE variable applies to each row named in the observation.
- The type of a column is identified when an observation in which the column name but no row name appears has the type *LOWERBD*, *UPPERBD*, *UNRSTRCT*, *FIXED*, *BINARY*, *INTEGER*, *BASIC*, *RHS*, *RHSSEN*, or *RANGE*. A column type can also be identified in an observation in which both column and row names appear and the row name has one of the preceding types.

- Each column is assumed to be a structural column in the model unless the column is identified as a right-hand-side vector, a right-hand-side change vector, or a range vector. A column can be identified as one of these types using either the keywords *RHS*, *RHSSEN*, or *RANGE* in the TYPE variable; the special column names _RHS_, _RHSSEN_, or _RANGE_; or the RHS, RHSSEN, or RANGE statements following the PROC LP statement.
- A TYPE variable beginning with the character * causes the observation to be interpreted as a comment.

When the column names appear in the Variable Summary in the PROC LP output, they are listed in alphabetical order. The row names appear in the order in which they appear in the problem data set.

Converting MPS Format

MPS input format was introduced by IBM. It has been a way of creating inputs for linear and integer programs. SASMPSXS is a SAS marco function that converts the standard MPS format to the sparse format of the LP procedure. The following is an example of the MPS format:

NAM	E	EXAMPLE			
* T	HIS IS DA	ATA FOR THE	PRODUCT MIX	PROBLEM.	
ROW	5				
N	PROFIT				
L	STAMP				
L	ASSEMB				
L	FINISH				
N	CHNROW				
N	PRICE				
COL	UMNS				
	DESK	STAMP	3.0000	0 ASSEMB	10.00000
	DESK	FINISH	10.0000	0 PROFIT	95.00000
	DESK	PRICE	175.0000	0	
	CHAIR	STAMP	1.5000	0 ASSEMB	6.00000
	CHAIR	FINISH	8.0000	0 PROFIT	41.00000
	CHAIR	PRICE	95.0000	0	
	CABINET	STAMP	2.0000	0 ASSEMB	8.00000
	CABINET	FINISH	8.0000	0 PROFIT	84.00000
	CABINET	PRICE	145.0000	0	
	BOOKCSE	STAMP	2.0000	0 ASSEMB	7.00000
	BOOKCSE	FINISH	7.0000	0 PROFIT	76.00000
	BOOKCSE	PRICE	130.0000	0 CHNROW	1.00000
RHS					
	TIME	STAMP	800.0000	0 ASSEMB	1200.0000
	TIME	FINISH	800.0000	0	
RAN	GES				
	т1	ASSEMB	900.0000	0	
BOU	NDS				
	UP	CHAIR	75.0000	0	
	LO	BOOKCSE	50.0000	0	
END	ATA				

In this example, the company tries to find an optimal product mix of four items: a DESK, a CHAIR, a CABINET, and a BOOKCASE. Each item is processed in a stamping department (STAMP), an assembly department (ASSEMB), and a finishing department (FINISH). The time each item requires in each department is given in the input data. Because of resource limitations, each department has an upper limit on the time available for processing. Furthermore, because of labor constraints, the assembly department must work at least 300 hours. Finally, marketing tells you not to make more than 75 chairs, to make at least 50 bookcases, and to find the range over which the selling price of a bookcase can vary without changing the optimal product mix.

The SASMPSXS macro function uses MPSFILE='FILENAME' as an argument to take an MPS input file. It then converts the file and saves the conversion to a default SAS data set, PROB. The FILENAME should include the path.

Running the following statements on the preceding example

```
%sasmpsxs(mpsfile='filename');
proc print data=prob;
run;
```

produces the sparse input form of the LP procedure:

OBS	_TYPE_	_COL_	_ROW1_	_COEF1_	_ROW2_	_COEF2	_
1	*OW			•		•	
2	FREE		PROFIT	•		•	
3	LE		STAMP	•		•	
4	LE		ASSEMB	•		•	
5	LE		FINISH	•		•	
6	FREE		CHNROW	•		•	
7	FREE		PRICE	•		•	
8	*OL	MNS		•		•	
9		DESK	STAMP	3.0	ASSEMB	10	
10		DESK	FINISH	10.0	PROFIT	95	
11		DESK	PRICE	175.0		•	
12		CHAIR	STAMP	1.5	ASSEMB	6	
13		CHAIR	FINISH	8.0	PROFIT	41	
14		CHAIR	PRICE	95.0		•	
15		CABINET	STAMP	2.0	ASSEMB	8	
16		CABINET	FINISH	8.0	PROFIT	84	
17		CABINET	PRICE	145.0		•	
18		BOOKCSE	STAMP	2	ASSEMB	7	
19		BOOKCSE	FINISH	7	PROFIT	76	
20		BOOKCSE	PRICE	130	CHNROW	1	
21	*HS			•		•	
22	RHS	TIME	STAMP	800	ASSEMB	1200	
23	RHS	TIME	FINISH	800		•	
24	*AN	ES		•		•	
25	RANGE	T1	ASSEMB	900		•	
26	*OU	DS		•		•	

27	UPPERBDD	CHAIR	UP	75	•
28	LOWERBDD	BOOKCSE	LO	50	•

SASMPSXS recognizes four MPS row types: *E*, *L*, *G* and *N*. It converts them into types *EQ*, *LE*, *GE* and *FREE*. Since objective rows, price change rows and free rows all share the same type *N* in the MPS format, you need a DATA step to assign proper types to the objective rows and price change rows.

```
data;
   set prob;
   if _type_='free' and _row1_='profit' then _type_='max';
   if _type_='free' and _row1_='chnrow' then _type_='pricesen';
run;
proc lp sparsedata;
run;
```

In the MPS format, the variable types include *LO*, *UP*, *FX*, *FR*, *MI*, and *BV*. The SASMPSXS macro converts them into types *LOWERBD*, *UPPERBD*, *FIXED*, *UN-RESTRICTED*, *-INFINITY*, and *BINARY*, respectively. Occasionally, you may define your own variable types, in which case, you need to add corresponding type handling entries in SASMPSXS.SAS program and use the SAS %INCLUDE macro to include the file at the beginning of your program. The SASMPSXS macro function can be found in the SAS sample library. The information on the MPS format can be obtained from Murtagh (1981).

SASMPSXS can take no arguments, or it can take one or two arguments. If no arguments are present, SASMPSXS assumes that the MPS input file has been saved to a SAS data set named RAW. The macro then takes information from that data set and converts it into the sparse form of the LP procedure. The RAW data set should have the following six variables:

```
data RAW;
    infile ...;
    input field1 $ 2-3 field2 $ 5-12
        field3 $ 15-22 field4 25-36
        field5 $ 40-47 field6 50-61;
    ...
run;
```

If the preceding MPS input data set has a name other than RAW, you can use MPSDATA=*SAS-data-set* as an argument in the SASMPSXS macro function. If you want the converted sparse form data set to have a name other than PROB, you can use LPDATA=*SAS-data-set* as an argument. The order of the arguments in the SASMPSXS macro function is not important.

The Reduced Costs, Dual Activities, and Current Tableau

The evaluation of reduced costs and the dual activities is independent of problem structure. For a basic solution, let B be the matrix composed of the basic columns of A and let N be the matrix composed of the nonbasic columns of A. The reduced cost associated with the i th variable is

$$(c^T - c^T_B B^{-1} A)_i$$

and the dual activity of the j th row is

$$(c_B^T B^{-1})_j$$

The Current Tableau is a section displayed when you specify either the TABLEAUPRINT option in the PROC LP statement or you specify the TABLEAU option in the PRINT statement. The output contains a row for each basic variable and a column for each nonbasic variable. In addition, there is a row for the reduced costs and a column for the product

 $B^{-1}b$

This column is labeled INV(B)*R. The body of the tableau contains the matrix

 $B^{-1}N$

Macro Variable _ORLP_

The LP procedure defines a macro variable named _ORLP_. This variable contains a character string that indicates the status of the procedure. It is set whenever the user gets control, at break points, and at procedure termination. The form of the _ORLP_ character string is STATUS= PHASE= OBJECTIVE= P_FEAS= D_FEAS= INT_ITER= INT_FEAS= ACTIVE= INT_BEST= PHASE1_ITER= PHASE2_ITER= PHASE3_ITER=. The terms are interpreted as follows:

STATUS=	the status of the current solution
PHASE=	the phase the procedure is in (1, 2, or 3)
OBJECTIVE=	the current objective value
P_FEAS=	whether the current solution is primal feasible
D_FEAS=	whether the current solution is dual feasible
INT_ITER=	the number of integer iterations performed
INT_FEAS=	the number of integer feasible solutions found
ACTIVE=	the number of active nodes in the current branch and bound tree
INT_BEST=	the best integer objective value found

PHASE1_ITER=	the number of iterations performed in phase 1
PHASE2_ITER=	the number of iterations performed in phase 2
PHASE3_ITER=	the number of iterations performed in phase 3

Table 3.5 shows the possible values for the nonnumeric terms in the string.

Table 3.5. Possible Values for Nonnumeric Terms

STATUS	P_FEASIBLE	D_FEASIBLE
SUCCESSFUL	YES	YES
UNBOUNDED	NO	NO
INFEASIBLE		
MAX_TIME		
MAX_ITER		
PIVOT		
BREAK		
INT_FEASIBLE		
INT_INFEASIBLE		
INT_MAX_ITER		
PAUSE		
FEASIBLEPAUSE		
IPAUSE		
PROXIMITYPAUSE		
ACTIVE		
RELAXED		
FATHOMED		
IPIVOT		
UNSTABLE		
SINGULAR		
MEMORY_ERROR		
IO_ERROR		
SYNTAX_ERROR		
SEMANTIC_ERROR		
BADDATA_ERROR		
UNKNOWN_ERROR		

This information can be used when PROC LP is one step in a larger program that needs to identify how the LP procedure terminated. Because _ORLP_ is a standard SAS macro variable, it can be used in the ways that all macro variables can be used (see the *SAS Guide to Macro Processing*).

Pricing

PROC LP performs multiple pricing when determining which variable will enter the basis at each pivot (Greenberg 1978). This heuristic can shorten execution time in many problems. The specifics of the multiple pricing algorithm depend on the value of the PRICETYPE= option. However, in general, when some form of multiple pricing is used, during the first iteration PROC LP places the PRICE= nonbasic columns

yielding the greatest marginal improvement to the objective function in a candidate list. This list identifies a subproblem of the original. On subsequent iterations, only the reduced costs for the nonbasic variables in the candidate list are calculated. This accounts for the potential time savings. When either the candidate list is empty or the subproblem is optimal, a new candidate list must be identified and the process repeats. Because identification of the subproblem requires pricing the complete problem, an iteration in which this occurs is called a *major iteration*. A *minor iteration* is an iteration in which only the subproblem is to be priced.

The value of the PRICETYPE= option determines the type of multiple pricing that is to be used. The types of multiple pricing include partial suboptimization (PRICE-TYPE=*PARTIAL*), complete suboptimization (PRICETYPE=*COMPLETE*), and complete suboptimization with dynamically varying the value of PRICE= option (PRICE-TYPE=*DYNAMIC*).

When partial suboptimization is used, in each minor iteration the nonbasic column in the subproblem yielding the greatest marginal improvement to the objective is brought into the basis and removed from the candidate list. The candidate list now has one less entry. At each subsequent iteration, another column from the subproblem is brought into the basis and removed from the candidate list. When there are either no remaining candidates or the remaining candidates do not improve the objective, the subproblem is abandoned and a major iteration is performed. If the objective cannot be improved on a major iteration, the current solution is optimal and PROC LP terminates.

Complete suboptimization is identical to partial suboptimization with one exception. When a nonbasic column from the subproblem is brought into the basis, it is replaced in the candidate list by the basic column that is leaving the basis. As a result, the candidate list does not diminish at each iteration.

When PRICETYPE=DYNAMIC, complete suboptimization is performed, but the value of the PRICE= option changes so that the ratio of minor to major iterations is within two units of the PRICE= option.

These heuristics can shorten execution time for small values of the PRICE= option. Care should be exercised in choosing a value from the PRICE= option because too large a value can use more time than if pricing were not used.

Scaling

Based on the SCALE= option specified, the procedure scales the coefficients of both the constraints and objective rows before iterating. This technique can improve the numerical stability of an ill-conditioned problem. If you want to modify the default matrix scaling used, which is SCALE=*NONE*, use the SCALE=*COLUMN*, SCALE=*ROW*, or SCALE=*BOTH* option in the PROC LP statement. If SCALE=*BOTH* is chosen, the matrix coefficients are scaled so that the largest element in absolute value in each row or column equals 1. They are scaled by columns first and then by rows. If SCALE=*COLUMN* (*ROW*), the matrix coefficients are scaled so that the largest element in absolute value in each column (row) equals 1.

Preprocessing

With the preprocessing option, you can identify redundant and infeasible constraints, improve lower and upper bounds of variables, fix variable values and improve coefficients and RHS values before solving a problem. Preprocessing can be applied to LP, IP and MIP problems. For an LP problem, it may significantly reduce the problem size. For an IP or MIP problem, it can often reduce the gap between the optimal solution and the solution of the relaxed problem, which could lead to a smaller search tree in the branch and bound algorithm. As a result, the CPU time may be reduced on many problems. Although there is no guarantee that preprocessing will always yield a faster solution, it does provide a highly effective approach to solving large and difficult problems.

Preprocessing is especially useful when the original problem causes numerical difficulties to PROC LP. Since preprocessing could identify redundant constraints and tighten lower and upper bounds of variables, the reformulated problem may eliminate the numerical difficulties in practice.

When a constraint is identified as redundant, its type is marked as *FREE* in the Constraint Summary. If a variable is fixed, its type is marked as *FIXED* in the Variables Summary. If a constraint is identified as infeasible, PROC LP stops immediately and displays the constraint name in the SAS log file. This capability sometimes gives valuable insight into the model or the formulation and helps establish if the model is reasonable and the formulation is correct.

For a large and dense problem, preprocessing may take a longer time for each iteration. To limit the number of preprocessings, use the PMAXIT= option. To stop any further preprocessings during the preprocessing stage, press the CTRL-BREAK key. PROC LP will enter phase 1 at the end of the current iteration.

Integer Programming

Formulations of mathematical programs often require that some of the decision variables take only integer values. Consider the formulation

 $\min c^T x$

subject to

$$Ax \ge =, =, \le b$$

 $\ell \le x \le u$
 $x_i \in S$ is integer

The set of indices S identifies those variables that must take only integer values. When S does not contain all of the integers between 1 and n, inclusive, problem (*mip*) is called a mixed-integer program. Otherwise, it is known as an integer program. Let $x^{opt}(mip)$ denote an optimal solution to (*mip*). An integer variable with bounds between 0 and 1 is also called a binary variable.

Specifying the Problem

An integer or mixed-integer problem can be solved with PROC LP. To solve this problem, you must identify the integer variables. You can do this with a row in the input data set that has the keyword *INTEGER* for the type variable. Any variable that has a nonmissing and nonzero value for this row is interpreted as an integer variable. It is important to note that integer variables must have upper bounds explicitly defined using the *UPPERBD* keyword. The values in the *INTEGER* row not only identify those variables that must be integral, but they also give an ordering to the integer variables that can be used in the solution technique.

You can follow the same steps to identify binary variables. For the binary variables, there is no need to supply any upper bounds.

Following the rules of sparse data input format, you can also identify individual integer or binary variables.

The Branch and Bound Technique

The branch and bound approach is used to solve integer and mixed-integer problems. The following discussion outlines the approach and explains how to use several options to control the procedure.

The branch and bound technique solves an integer program by solving a sequence of linear programs. The sequence can be represented by a tree, with each node in the tree being identified with a linear program that is derived from the problems on the path leading to the root of the tree. The root of the tree is identified with a linear program that is identical to (mip) except that S is empty. This relaxed version of (mip), lp(0), can be written as:

$$x^{opt}(0) = \min c^T x$$

subject to

$$Ax \ge , =, \le b$$
$$\ell \le x \le u$$

The branch and bound approach generates linear programs along the nodes of the tree using the following scheme. Consider $x^{opt}(0)$, the optimal solution to lp(0). If $x^{opt}(0)_i$ is integer for all $i \in S$, then $x^{opt}(0)$ is optimal in (mip). Suppose for some $i \in S$, $x^{opt}(0)_i$ is nonintegral. In that case, define two new problems (lp(1)) and (lp(2)), descendents of the parent problem (lp(0)). The problem (lp(1)) is identical to (lp(0)) except for the additional constraint

$$x_i \le \lfloor x^{opt}(0)_i \rfloor$$

and the problem (lp(2)) is identical to (lp(0)) except for the additional constraint

$$x_i \ge \lceil x^{opt}(0)_i \rceil$$

The notation $\lceil y \rceil$ means the smallest integer greater than or equal to y, and the notation $\lfloor y \rfloor$ means the largest integer less than or equal to y. Note that the two new problems do not have $x^{opt}(0)$ as a feasible solution, but because the solution to (mip) must satisfy one of the preceding constraints, $x^{opt}(mip)_i$ must satisfy one of the new constraints. The two problems thus defined are called active nodes in the branch and bound tree, and the variable i is called the branching variable.

Next, the algorithm chooses one of the problems associated with an active node and attempts to solve it using the dual simplex algorithm. The problem may be infeasible, in which case the problem is dropped. If it can be solved, and it in turn does not have an integer solution (that is a solution for which x_i is integer for all $i \in S$), then it defines two new problems. These new problems each contain all of the constraints of the parent problems plus the appropriate additional one.

Branching continues in this manner until either there are no active nodes or an integer solution is found. When an integer solution is found, its objective value provides a bound for the objective of (mip). In particular, if z is the objective value of the current best integer solution, then any active problems whose parent problem has objective value $\geq z$ can be discarded (assuming that the problem is a minimization). This can be done because all problems that descend from this parent will also have objective value $\geq z$. This technique is known as fathoming. When there are no active nodes remaining to be solved, the current integer solution is optimal in (mip). If no integer solution has been found, then (mip) is (integer) infeasible.

It is important to realize that integer programs are NP-complete. Roughly speaking, this means that the effort required to solve them grows exponentially with the size of the problem. For example, a problem with 10 binary variables can, in the worst case, generate $2^{10} = 1024$ nodes in the branch and bound tree. A problem with 20 binary variables can, in the worst case, generate $2^{20} = 1048576$ nodes in the branch and bound tree. Although the algorithm is unlikely to have to generate every single possible node, the need to explore even a small fraction of the potential number of nodes for a large problem can be resource intensive.

The Integer Iteration Log

To help monitor the growth of the branch and bound tree, the LP procedure reports on the status of each problem that is solved. The report, displayed in the Integer Iteration Log, can be used to reconstruct the branch and bound tree. Each row in the report describes the results of the attempted solution of the linear program at a node in the tree. In the following discussion, a problem on a given line in the log is called the current problem. The following columns are displayed in the report:

Iter	identifies the number of the branch and bound iteration.

Problem identifies how the current problem fits in the branch and bound tree.

Condition reports the result of the attempted solution of the current problem. Values for Condition are

ACTIVE	The current problem was solved successfully.
INFEASIBLE	The current problem is infeasible.

	FATHOMED	The current problem cannot lead to an improved integer solution and therefore it is dropped.	
	SINGULAR	A singular basis was encountered in at- tempting to solve the current problem. Solution of this relaxed problem is sus- pended and will be attempted later if necessary.	
	SUBOPTIMAL	The current problem has an integer fea- sible solution.	
Objective	reports the objective value of the current problem.		
Branched	names the variable that is branched in subtrees defined by the de- scendents of this problem.		
Value	gives the current value of the variable named in the column labeled Branched.		
Sinfeas	gives the sum of the integ to the current problem.	ger infeasibilities in the optimal solution	
Active	reports the total number of nodes currently active in the branch and bound tree.		
Proximity	reports the gap between the best integer solution and the current lower (upper for maximizations) bound of all active nodes.		

To reconstruct the branch and bound tree from this report, consider the interpretation of iteration j. If Iter=j and Problem=k, then the problem solved on iteration j is identical to the problem solved on iteration |k| with an additional constraint. If k > 0, then the constraint is an upper bound on the variable named in the Branched column on iteration |k|. On the other hand, if k < 0 then the constraint is a lower bound on that variable. The value of the bound can be obtained from the value of Value in iteration |k| as described in the previous section.

Example 3.8 in the section "Examples" on page 123 shows an Integer Iteration Log in its output.

Controlling the Branch and Bound Search

There are several options you can use to control branching. This is accomplished by controlling the program's choice of the branching variable and of the next active node. In the discussion that follows, let

$$f_i(k) = x^{opt}(k)_i - \lfloor x^{opt}(k)_i \rfloor$$

where $x^{opt}(k)$ is the optimal solution to the problem solved on iteration k.

The CANSELECT= option directs the choice of the next active node. Valid keywords for this option include *LIFO*, *FIFO*, *OBJ*, *PROJECT*, *PSEUDOC*, and *ERROR*. The following list describes the action that each of these causes when the procedure must choose for solution a problem from the list of active nodes.

- *LIFO* chooses the last problem added to the tree of active nodes. This search has the effect of a depth first search of the branch and bound tree.
- *FIFO* chooses the first node added to the tree of active nodes. This search has the effect of a breadth first search of the branch and bound tree.
- *OBJ* chooses the problem whose parent has the least (largest if the problem is a maximization) objective value.
- **PROJECT** chooses the problem with the largest (least if the problem is a maximization) projected objective value. The projected objective value is evaluated using the sum of integer infeasibilities, s(k), associated with an active node (lp(k)), defined by:

$$s(k) = \sum_{i \in \mathcal{S}} \min\{f_i(k), (1 - f_i(k))\}.$$

An empirical measure of the rate of increase (decrease) in the objective value is defined as

$$\lambda = (z^* - z(0))/s(0)$$

where

z(k) is the optimal objective value for (lp(k)).

 z^* is the objective value of the current best integer solution.

The projected objective value for problems (lp(k+1)) and (lp(k+2)) is defined as

$$z(k) + \lambda s(k)$$

PSEUDOC chooses the problem with the largest (least if the problem is a maximization) projected pseudocost. The projected pseudocost is evaluated using the weighted sum of infeasibilities $s_w(k)$ associated with an active problem (lp(k)), defined by

$$s_w(k) = \sum_{i \in \mathcal{S}} \min\{d_i(k) f_i(k), u_i(k)(1 - f_i(k))\}$$

The weights u_i and d_i are initially equal to the absolute value of the *i*th objective coefficient and are updated at each integer iteration. They are modified by examining the empirical marginal change in the objective as additional constraints are placed on the variables in S along the path from (lp(0)) to a node associated with an integer feasible solution. In particular, if the definition of problems (lp(k+1)) and (lp(k+2)) from parent (lp(k)) involve the addition of constraints $x_i \leq \lfloor x^{opt}(k)_i \rfloor$ and $x_i \geq \lceil x^{opt}(k)_i \rceil$, respectively, and one of them is on a path to an integer feasible solution, then only one of the following is true:

$$d_i(k) = (z(k+1) - z(k))/f_i(k)$$

$$u_i(k) = (z(k+2) - z(k))/(1 - f_i(k))$$

Note the similarity between $s_w(k)$ and s(k). The weighted quantity $s_w(k)$ accounts to some extent for the influence of the objective function. The projected pseudocost for problems (lp(k + 1)) and (lp(k + 2)) is defined as

$$z_w(k) \equiv z(k) + s_w(k)$$

ERROR chooses the problem with the largest error. The error associated with problems (lp(k + 1)) and (lp(k + 2)) is defined as

$$(z^* - z_w(k))/(z^* - z(k))$$

The BACKTRACK= option controls the search for the next problem. This option can take the same values as the CANSELECT= option. In addition to the case outlined under the DELTAIT= option, backtracking is required as follows based on the CANSELECT= option in effect:

- If CANSELECT=*LIFO* and there is no active node in the portion of the active tree currently under exploration with a bound better than the value of WOB-JECTIVE=, then the procedure must backtrack.
- If CANSELECT=*FIFO*, *PROJECT*, *PSEUDOC*, and *ERROR* and the bound corresponding to the node under consideration is not better than the value of WOBJECTIVE=, then the procedure must backtrack.

The default value is *OBJ*.

The VARSELECT= option directs the choice of branching variable. Valid keywords for this option include *CLOSE*, *FAR*, *PRIOR*, *PSEUDOC*, *PRICE*, and *PENALTY*. The following list describes the action that each of these causes when $x^{opt}(k)$, an optimal solution of problem (lp(k)), is used to define active problems (lp(k + 1)) and (lp(k + 2)).

CLOSE	chooses as branching variable the variable x_i such that i minimizes
	$\{\min\{f_i(k), (1 - f_i(k))\} \mid i \in \mathcal{S} ext{ and IEPSILON} \leq f_i(k) \leq 1 - ext{IEPSILON} \}$
FAR	chooses as branching variable the variable x_i such that i maximizes
	$\{\min\{f_i(k),(1-f_i(k))\} \mid i \in \mathcal{S} ext{ and IEPSILON} \leq f_i(k) \leq 1 - ext{IEPSILON} \}$
PRIOR	chooses as branching variable x_i such that $i \in S$, $x^{opt}(k)_i$ is non- integral, and variable x_i has the minimum value in the <i>INTEGER</i> row in the input data set. This choice for the VARSELECT= option is recommended when you have enough insight into the model to identify those integer variables that have the most significant effect on the objective value.

PENALTY chooses as branching variable x_i such that $i \in S$ and a bound on the increase in the objective of (lp(k)) (penalty) resulting from adding the constraint

$$x_i \leq \lfloor x^{opt}(k)_i \rfloor$$
 or
 $x_i \geq \lceil x^{opt}(k)_i \rceil$

is maximized. The bound is calculated without pivoting using techniques of sensitivity analysis (Garfinkel and Nemhauser 1972). Because the cost of calculating the maximum penalty can be large if S is large, you may want to limit the number of variables in S for which the penalty is calculated. The penalty is calculated for PENALTYDEPTH= variables in S.

PRICE chooses as branching variable x_i such that $i \in S$, $x^{opt}(k)_i$ is nonintegral, and variable x_i has the minimum price coefficient (maximum for maximization).

PSEUDOC chooses as branching variable the variable x_i such that *i* maximizes

$$\{ \min\{d_i f_i(k), u_i(1 - f_i(k))\} \mid i \in S \text{ and IEPSILON} \\ \leq f_i(k) \leq 1 - \text{IEPSILON} \}$$

The weights u_i and d_i are initially equal to the absolute value of the *i*th objective coefficient and are updated whenever an integer feasible solution is encountered. See the discussion on the CANS-ELECT= option for details on the method of updating the weights.

Customizing Search Heuristics

Often a good heuristic for searching the branch and bound tree of a problem can be found. You are tempted to continue using this heuristic when the problem data changes but the problem structure remains constant. The ability to reset procedure options interactively enables you to experiment with search techniques in an attempt to identify approaches that perform well. Then you can easily reapply these techniques to subsequent problems.

For example, the PIP branch and bound strategy (Crowder, Johnson, and Padberg 1983) describes one such heuristic. The following program uses a similar strategy. Here, the *OBJ* rule (choose the active node with least parent objective function in the case of a minimization problem) is used for selecting the next active node to be solved until an integer feasible solution is found. Once such a solution is found, the search procedure is changed to the *LIFO* rule: choose the problem most recently placed in the list of active nodes.

```
proc lp canselect=obj ifeasiblepause=1;
run;
    reset canselect=lifo ifeasiblepause=9999999;
run;
```

Further Discussions on AUTO and CONTROL= options

Consider a minimization problem. At each integer iteration, PROC LP will select a node to solve from a pool of active nodes. The best bound strategy (CANSE-LECT=OBJ) will pick the node with the smallest projected objective value. This strategy improves the lower bound of the integer program and usually takes fewer integer iterations. One disadvantage is that PROC LP must recalculate the inverse of the basis matrix at almost every integer iteration; such recalculation is relatively expensive. Another disadvantage is that this strategy does not pay attention to improving the upper bound of the integer program. Thus the number of active nodes tends to grow rapidly if PROC LP cannot quickly find an optimal integer solution.

On the other hand, the LIFO strategy is very efficient and does not need to calculate the inverse of the basis matrix unless the previous node is fathomed. It is a depth first strategy so it tends to find an integer feasible solution quickly. However, this strategy will pick nodes locally and usually will take longer integer iterations than the best bound strategy.

There is another strategy that is often overlooked. Here it is called the **best upper bound** strategy. With this strategy, each time you select an active node each time, instead of picking the node with the smallest projected objective value, you select the one with the largest projected objective value. This strategy is as efficient as the LIFO strategy. Moreover, it selects active nodes globally. This strategy tries to improve the upper bound of the integer program by searching for new integer feasible solutions. It also fathoms active nodes quickly and keeps the total number of active nodes below current level. A disadvantage is that this strategy may evaluate more nodes that do not have any potential in finding an optimal integer solution.

The best bound strategy has the advantage of improving the lower bound. The LIFO strategy has the advantages of efficiency and finding a local integer feasible solution. The best upper bound strategy has the advantages of keeping the size of active nodes under control and at the same time trying to identify any potential integer feasible solution globally.

Although the best bound strategy is generally preferred, in some instances other strategies may be more effective. For example, if you have found an integer optimal solution but you do not know it, you still have to enumerate all possible active nodes. Then the three strategies will basically take the same number of integer iterations after an optimal solution is found but not yet identified. Since the LIFO and best upper bound strategies are very efficient per integer iteration, both will outperform the best bound strategy.

Since no one strategy suits all situations, a hybrid strategy has been developed to increase applicability. The CONTROL= option combines the above three strategies naturally and provides a simple control parameter in [0, 1] dealing with different integer programming problems and different solution situations. The AUTO option automatically sets and adjusts the CONTROL= parameter so that you do not need to know any problem structure or decide a node selection strategy in advance.

Since the LIFO strategy is less costly, you should use it as much as possible in the combinations. The following process is called a **diving process**. Starting from an

active node, apply the LIFO strategy as much as you can until the current node becomes feasible or is fathomed, or exceeds a preset limit. During this process, there is no inverse matrix calculation involved except for the first node. When the diving process is over, apply one of the three strategies to select the next starting node. One set of combinations is called a cycle.

The control parameter r controls the frequency of the three strategies being applied and the depth of the diving process in a cycle. It starts with a pure best bound strategy at r = 0, and then gradually increases the frequency of the diving processes and their depths as r increases. At r = 0.5, one cycle contains a best bound strategy plus a full diving process. After r = 0.5, the number of the diving processes will gradually increase in a cycle. In addition, the best upper bound strategy will join the cycle. As r continues to increase, the frequency of the best upper bound strategy will increase. At r = 1, it becomes a pure best upper bound strategy.

The AUTO option will automatically adjust the value of the CONTROL= option. At the start, it sets CONTROL=0.7, which emphasizes the finding of an upper bound. After an integer feasible solution is found, it sets CONTROL=0.5, which emphasizes efficiency and lower bound improvement. When the number of active nodes grows over the default or user defined limit m, the number indicates that a better upper bound is needed. The AUTO option will start to increase the value of CONTROL= from 0.5. If the size of the active nodes continues to grow, so will the value of CONTROL= option. When the size of active nodes reaches to the default or user defined limit n, CONTROL= will be set to 1. At this moment, the growth of active nodes is stopped. When the size of active nodes reduces, AUTO will decrease the value of CONTROL= option.

You can use other strategies to improve the lower bound by setting CANSELECT= to other options.

Saving and Restoring the List of Active Nodes

The list of active nodes can be saved in a SAS data set for use at a subsequent invocation of PROC LP. The ACTIVEOUT= option in the PROC LP statement names the data set into which the current list of active nodes is saved when the procedure terminates due to an error termination condition. Examples of such conditions are time limit exceeded, integer iterations exceeded, and phase 3 iterations exceeded. The AC-TIVEIN= option in the PROC LP statement names a data set that can be used to initialize the list of active nodes. To achieve the greatest benefit when restarting PROC LP, use the PRIMALOUT= and PRIMALIN= options in conjunction with the AC-TIVEOUT= and ACTIVEIN= options. See Example 3.10 in the section "Examples" on page 123 for an illustration.

Sensitivity Analysis

Sensitivity analysis is a technique for examining the effects of changes in model parameters on the optimal solution. The analysis enables you to examine the size of a perturbation to the right-hand-side or objective vector by an arbitrary change vector for which the basis of the current optimal solution remains optimal.

Note: When sensitivity analysis is performed on integer constrained problems, the integer variables are fixed at the value they obtained in the integer optimal solution. Therefore, care must be used when interpreting the results of such analyses. Care must also be taken when preprocessing is enabled, because preprocessing usually alters the original formulation.

Right-Hand-Side Sensitivity Analysis

Consider the problem $(lpr(\phi))$

$$x^{opt}(\phi) = \min c^T x$$

subject to

$$Ax \ge , =, \le b + \phi r$$

 $\ell \le x \le u$

where r is a right-hand-side change vector.

Let $x^{opt}(\phi)$ denote an optimal basic feasible solution to $(lpr(\phi))$. PROC LP can be used to examine the effects of changes in ϕ on the solution $x^{opt}(0)$ of problem (lpr(0)). For the basic solution $x^{opt}(0)$, let B be the matrix composed of the basic columns of A and let N be the matrix composed of the nonbasic columns of A. For the basis matrix B, the basic components of $x^{opt}(0)$, written as $x^{opt}(0)_B$, can be expressed as

$$x^{opt}(0)_B = B^{-1}(b - Nx^{opt}(0)_N)$$

Furthermore, because $x^{opt}(0)$ is feasible

$$\ell_B \le B^{-1}(b - Nx^{opt}(0)_N) \le u_B$$

where ℓ_B is a column vector of the lower bounds on the structural basic variables, and u_B is a column vector of the upper bounds on the structural basic variables. For each right-hand-side change vector r identified in the RHSSEN statement, PROC LP finds an interval $[\phi_{min}, \phi_{max}]$ such that

$$\ell_B \le B^{-1}(b + \phi r - Nx^{opt}(0)_N) \le u_B$$

for $\phi \in [\phi_{min}, \phi_{max}]$. Furthermore, because changes in the right-hand side do not affect the reduced costs, for $\phi \in [\phi_{min}, \phi_{max}]$

$$x^{opt}(\phi)^T = ((B^{-1}(b + \phi r - Nx^{opt}(0)_N))^T, x^{opt}(0)_N^T)$$

is optimal in $(lpr(\phi))$.

For $\phi = \phi_{min}$ and $\phi = \phi_{max}$, PROC LP reports the following:

- the names of the leaving variables
- the value of the optimal objective in the modified problems
- the the RHS values in the modified problems
- the solution status, reduced costs and activities in the modified problems

The leaving variable identifies the basic variable *i* that first reaches either the lower bound ℓ_i or the upper bound u_i as ϕ reaches ϕ_{min} or ϕ_{max} . This is the basic variable that would leave the basis to maintain primal feasibility. Multiple RHSSEN variables can appear in a problem data set.

Price Sensitivity Analysis

Consider the problem $(lpp(\phi))$

$$x^{opt}(\phi) = \min \left(c + \phi r \right)^T x$$

subject to

$$egin{array}{l} Ax\geq,=,\leq b\ \ell\leq x\leq u \end{array}$$

where r is a price change vector.

Let $x^{opt}(\phi)$ denote an optimal basic feasible solution to $(lpp(\phi))$. PROC LP can be used to examine the effects of changes in ϕ on the solution $x^{opt}(0)$ of problem (lpp(0)). For the basic solution $x^{opt}(0)$, let B be the matrix composed of the basic columns of A and let N be the matrix composed of the nonbasic columns of A. For basis matrix B, the reduced cost associated with the *i* th variable can be written as

$$rc_i(\phi) = ((c + \phi r)_N^T - (c + \phi r)_B^T B^{-1} N)_i$$

where $(c + \phi r)_N$ and $(c + \phi r)_B$ is a partition of the vector of price coefficients into nonbasic and basic components. Because $x^{opt}(0)$ is optimal in $(lpp \ (0))$, the reduced costs satisfy

$$rc_i(\phi) \ge 0$$

if the nonbasic variable in column i is at its lower bound and

$$rc_i(\phi) \leq 0$$

if the nonbasic variable in column i is at its upper bound.

For each price coefficient change vector r identified with the keyword PRICESEN in the TYPE variable, PROC LP finds an interval $[\phi_{min}, \phi_{max}]$ such that for $\phi \in [\phi_{min}, \phi_{max}]$

$$rc_i(\phi) \ge 0$$

if the nonbasic variable in column i is at its lower bound and

$$rc_i(\phi) \leq 0$$

if the nonbasic variable in column *i* is at its upper bound. Because changes in the price coefficients do not affect feasibility, for $\phi \in [\phi_{min}, \phi_{max}]$, $x^{opt}(\phi)$ is optimal in $(lpp(\phi))$. For $\phi = \phi_{min}$ and $\phi = \phi_{max}$, PROC LP reports the following:

- the names of entering variables
- the value of the optimal objective in the modified problems
- the price coefficients in the modified problems
- the solution status, reduced costs, and activities in the modified problems

The entering variable identifies the variable whose reduced cost first goes to zero as ϕ reaches ϕ_{min} or ϕ_{max} . This is the nonbasic variable that would enter the basis to maintain optimality (dual feasibility). Multiple PRICESEN variables may appear in a problem data set.

Range Analysis

Range analysis is sensitivity analysis for specific change vectors. As with the sensitivity analysis case, care must be used in interpreting the results of range analysis when the problem has integers or the preprocessing option is enabled.

Right-Hand-Side Range Analysis

The effects on the optimal solution of changes in each right-hand-side value can be studied using the RANGERHS option in the PROC LP or RESET statement. This option results in sensitivity analysis for the m right-hand-side change vectors specified by the columns of the $m \times m$ identity matrix.

Price Range Analysis

The effects on the optimal solution of changes in each price coefficient can be studied using the RANGEPRICE option in the PROC LP or RESET statement. This option results in sensitivity analysis for the n price change vectors specified by the rows of the $n \times n$ identity matrix.

Parametric Programming

Sensitivity analysis and range analysis examine how the optimal solution behaves with respect to perturbations of model parameter values. These approaches assume that the basis at optimality is not allowed to change. When greater flexibility is desired and a change of basis is acceptable, parametric programming can be used. As with the sensitivity analysis case, care must be used in interpreting the results of parametric programming when problem has integers or preprocessing option is enabled.

Right-Hand-Side Parametric Programming

As discussed in the "Right-Hand-Side Sensitivity Analysis" section, for each righthand-side change vector r, PROC LP finds an interval $[\phi_{min}, \phi_{max}]$ such that for $\phi \in [\phi_{min}, \phi_{max}]$

 $x^{opt}(\phi)^T = ((B^{-1}(b + \phi r - Nx^{opt}(0)_N))^T, x^{opt}(0)_N^T)$

is optimal in $(lpr(\phi))$ for the fixed basis *B*. Leaving variables that inhibit further changes in ϕ without a change in the basis *B* are associated with the quantities ϕ_{min} and ϕ_{max} . By specifying RHSPHI= Φ in either the PROC LP statement or in the RESET statement, you can examine the solution $x^{opt}(\phi)$ as ϕ increases or decreases from 0 to Φ .

When RHSPHI= Φ is specified, the procedure first finds the interval $[\phi_{min}, \phi_{max}]$ as described previously. Then, if $\phi \in [\phi_{min}, \phi_{max}]$, no further investigation is needed. However, if $\Phi > \phi_{max}$ or $\Phi < \phi_{min}$, then the procedure attempts to solve the new problem $(lpr(\Phi))$. To accomplish this, it pivots the leaving variable out of the basis while maintaining dual feasibility. If this new solution is primal feasible in $(lpr(\Phi))$, no further investigation is needed; otherwise, the procedure identifies the new leaving variable and pivots it out of the basis, again maintaining dual feasibility. Dual pivoting continues in this manner until a solution that is primal feasible in $(lpr(\Phi))$ is identified. Because dual feasibility is maintained at each pivot, the $(lpr(\Phi))$ primal feasible solution is optimal.

At each pivot, the procedure reports on the variables that enter and leave the basis, the current range of ϕ , and the objective value. When $x^{opt}(\Phi)$ is found, it is displayed. If you want the solution $x^{opt}(\phi)$ at each pivot, then specify the PARAPRINT option in either the PROC LP or the RESET statement.

Price Parametric Programming

As discussed in the "Price Sensitivity Analysis" section, for each price change vector r, PROC LP finds an interval $[\phi_{min}, \phi_{max}]$ such that for each $\phi \in [\phi_{min}, \phi_{max}]$

$$rc_i(\phi) = ((c + \phi r)_N^T - (c + \phi r)_B^T B^{-1} N)_i$$

satisfies the conditions for optimality in $(lpp(\phi))$ for the fixed basis *B*. Entering variables that inhibit further changes in ϕ without a change in the basis *B* are associated with the quantities ϕ_{min} and ϕ_{max} . By specifying PRICEPHI= Φ in either the PROC LP statement or the RESET statement, you can examine the solution $x^{opt}(\phi)$ as ϕ increases or decreases from 0 to Φ .

When PRICEPHI= Φ is specified, the procedure first finds the interval $[\phi_{min}, \phi_{max}]$, as described previously. Then, if $\Phi \in [\phi_{min}, \phi_{max}]$, no further investigation is needed. However, if $\Phi > \phi_{max}$ or $\Phi < \phi_{min}$, the procedure attempts to solve the new problem $(lpp(\Phi))$. To accomplish this, it pivots the entering variable into the basis while maintaining primal feasibility. If this new solution is dual feasible in $(lpp(\Phi))$, no further investigation is needed; otherwise, the procedure identifies the

112 • Chapter 3. The LP Procedure

new entering variable and pivots it into the basis, again maintaining primal feasibility. Pivoting continues in this manner until a solution that is dual feasible in $(lpp(\Phi))$ is identified. Because primal feasibility is maintained at each pivot, the $(lpp(\Phi))$ dual feasible solution is optimal.

At each pivot, the procedure reports on the variables that enter and leave the basis, the current range of ϕ , and the objective value. When $x^{opt}(\Phi)$ is found, it is displayed. If you want the solution $x^{opt}(\phi)$ at each pivot, then specify the PARAPRINT option in either the PROC LP or the RESET statement.

Interactive Facilities

The interactive features of the LP procedure enable you to examine intermediate results, perform sensitivity analysis, parametric programming, and range analysis, and control the solution process.

Controlling Interactive Features

You can gain control of the LP procedure for interactive processing by setting a break point or pressing the CTRL-BREAK key combination, or when certain error conditions are encountered:

- when a feasible solution is found
- at each pivot of the simplex algorithm
- when an integer feasible solution is found
- at each integer pivot of the branch and bound algorithm
- after the data are read but before iteration begins
- after at least one integer feasible solution has been found which is within desirable proximity of optimality
- after the problem has been solved but before results are displayed

When the LP procedure pauses, you can enter any of the interactive statements RE-SET, PIVOT, IPIVOT, PRINT, SHOW, QUIT, and RUN.

Break points are set using the FEASIBLEPAUSE, PAUSE, IFEASIBLEPAUSE=, IPAUSE=, PROXIMITYPAUSE=, READPAUSE, and ENDPAUSE options. The LP procedure displays a message on the SAS log when it gives you control because of encountering one of these break points.

During phase 1, 2, or 3, the CTRL-BREAK key pauses the LP procedure and releases the control at the beginning of next iteration.

The error conditions, which usually cause the LP procedure to pause, include time limit exceeded, phase 1 iterations exceeded, phase 2 iterations exceeded, phase 3 iterations exceeded, and integer iterations exceeded. You can use the RESET statement to reset the option that caused the error condition.

The PIVOT and IPIVOT statements result in control being returned to you after a single simplex algorithm pivot and an integer pivot. The PRINT and SHOW statements display current solution information and return control to you. On the other hand, the QUIT statement requests that you leave the LP procedure immediately. If you want to quit but save output data sets, then type QUIT/SAVE. The RUN statement requests the LP procedure to continue its execution immediately.

Displaying Intermediate Results

Once you have control of the procedure, you can examine the current values of the options and the status of the problem being solved using the SHOW statement. All displaying done by the SHOW statement goes to the SAS log.

Details about the current status of the solution are obtained using the PRINT statement. The various display options enable you to examine parts of the variable and constraint summaries, display the current tableau, perform sensitivity analysis on the current solution, and perform range analysis.

Interactive Facilities in Batch Mode

All of the interactive statements can be used when processing in batch mode. This is particularly convenient when the interactive facilities are used to combine different search strategies in solving integer problems.

Sensitivity Analysis

Two features that enhance the ability to perform sensitivity analysis need further explanation. When you specify /SENSITIVITY in a PRINT COLUMN(*colnames*) statement, the LP procedure defines a new change row to use in sensitivity analysis and parametric programming. This new change row has a +1 entry for each variable listed in the PRINT statement. This enables you to define new change rows interactively.

When you specify /SENSITIVITY in a PRINT ROW (*rownames*) statement, the LP procedure defines a new change column to use in sensitivity analysis and parametric programming. This new change column has a +1 entry for each right-hand-side coefficient listed in the PRINT statement. This enables you to define new change columns interactively.

In addition, you can interactively change the RHSPHI= and PRICEPHI= options using the RESET statement. This enables you to perform parametric programming interactively.

Memory Management

There are no restrictions on the problem size in the LP procedure. The number of constraints and variables in a problem that PROC LP can solve depends on the host platform, the available memory, and the available disk space for utility data sets.

Memory usage is affected by a great many factors including the density of the technological coefficient matrix, the model structure, and the density of the decomposed basis matrix. The algorithm requires that the decomposed basis fit completely in memory. Any additional memory is used for nonbasic columns. The partition between the decomposed basis and the nonbasic columns is dynamic so that as the inverse grows, which typically happens as iterations proceed, more memory is available to it and less is available for the nonbasic columns. The LP procedure determines the initial size of the decomposed basis matrix. If the area used is too small, PROC LP must spend time compressing this matrix, which degrades performance. If PROC LP must compress the decomposed basis matrix on the average more than 15 times per iteration, then the size of the memory devoted to the basis is increased. If the work area cannot be made large enough to invert the basis, an error return occurs. On the other hand, if PROC LP compresses the decomposed basis matrix on the average once every other iteration, then memory devoted to the decomposed basis is decreased, freeing memory for the nonbasic columns.

For many models, memory constraints are not a problem because both the decomposed basis and all the nonbasic columns will have no problem fitting. However, when the models become large relative to the available memory, the algorithm tries to adjust memory distribution in order to solve the problem. In the worst cases, only one nonbasic column fits in memory with the decomposed basis matrix.

Problems involving memory use can occur when solving mixed-integer problems. Data associated with each node in the branch and bound tree must be kept in memory. As the tree grows, competition for memory by the decomposed basis, the nonbasic columns, and the branch and bound tree may become critical. If the situation becomes critical, the procedure automatically switches to branching strategies that use less memory. However, it is possible to reach a point where no further processing is possible. In this case, PROC LP terminates on a memory error.

Output Data Sets

The LP procedure can optionally produce four output data sets. These are the AC-TIVEOUT, PRIMALOUT, DUALOUT, and TABLEAUOUT data sets. Each contains two variables that identify the particular problem in the input data set. These variables are

_OBJ_ID_	identifies the objective function ID.
_RHS_ID_	identifies the right-hand-side variable.

Additionally, each data set contains other variables, which are discussed below.

ACTIVEOUT= Data Set

The ACTIVEOUT= data set contains a representation of the current active branch and bound tree. You can use this data set to initialize the branch and bound tree to continue iterations on an incompletely solved problem. Each active node in the tree generates two observations in this data set. The first is a *LOWERBD* observation that is used to reconstruct the lower-bound constraints on the currently described active node. The second is an *UPPERBD* observation that is used to reconstruct the upperbound constraints on the currently described active node. In addition to these, an observation that describes the current best integer solution is included. The data set contains the following variables:

- _STATUS_ contains the keywords *LOWERBD*, *UPPERBD*, and *INTBEST* for identifying the type of observation.
- _PROB_ contains the problem number for the current observation.

- _OBJECT_ contains the objective value of the parent problem that generated the current observation's problem.
- _SINFEA_ contains the sum of the integer infeasibilities of the current observation's problem.
- _PROJEC_ contains the data needed for CANSELECT=*PROJECT* when the branch and bound tree is read using the ACTIVEIN= option.
- _PSEUDO_ contains the data needed for CANSELECT=*PSEUDOC* when the branch and bound tree is read using the ACTIVEIN= option.
- INTEGER VARIABLES integer constrained structural variables are also included in the ACTIVEOUT= data set. For each observation, these variables contain values for defining the active node in the branch and bound tree.

PRIMALOUT= Data Set

The PRIMALOUT= data set contains the current primal solution. If the problem has integer-constrained variables, the PRIMALOUT= data set contains the current best integer feasible solution. If none have been found, the PRIMALOUT= data set contains the relaxed solution. In addition to _OBJ_ID_ and _RHS_ID_, the data set contains the following variables:

- _VAR_ identifies the variable name.
- **_TYPE_** identifies the type of the variable as specified in the input data set. Artificial variables are labeled as type *ARTIFCL*.
- **_STATUS_** identifies whether the variable is basic, nonbasic, or at an upper bound in the current solution.
- _LBOUND_ contains the input lower bound on the variable unless the variable is integer-constrained and an integer solution is given. In this case, _LBOUND_ contains the lower bound on the variable needed to realize the integer solution on subsequent calls to PROC LP when using the PRIMALIN= option.
- _VALUE_ identifies the value of the variable in the current solution or the current best integer feasible solution.
- _UBOUND_ contains the input upper bound on the variable unless the variable is integer-constrained and an integer solution is given. In this case, _UBOUND_ contains the upper bound on the variable needed to realize the integer solution on subsequent calls to PROC LP when using the PRIMALIN= option.
- _PRICE_ contains the input price coefficient of the variable.
- _R_COST_ identifies the value of the reduced cost in the current solution. Example 3.3 in the "Examples" section shows a typical PRI-MALOUT= data set. Note that it is necessary to include the information on objective function and right-hand side in order to distinguish problems in multiple problem data sets.

DUALOUT= Data Set

The DUALOUT= data set contains the dual solution for the current solution. If the problem has integer-constrained variables, the DUALOUT= data set contains the dual for the current best integer solution, if any. Otherwise it contains the dual for the relaxed solution. In addition to _OBJ_ID_ and _RHS_ID_, it contains the following variables:

_ROW_ID_	identifies the row or constraint name.
TYPE	identifies the type of the row as specified in the input data set.
RHS	gives the value of the right-hand-side on input.
_L_RHS_	gives the lower bound for the row evaluated from the input right- hand-side value, the TYPE of the row, and the value of the RANGE variable for the row.
VALUE	gives the value of the row, at optimality, excluding logical variables.
_U_RHS_	gives the upper bound for the row evaluated from the input right- hand-side value, the TYPE of the row, and the value of the RANGE variable for the row.
DUAL	gives the value of the dual variable associated with the row.

TABLEAUOUT= Data Set

The TABLEAUOUT= data set contains the current tableau. Each observation, except for the first, corresponds to a basic variable in the solution. The observation labeled R_COSTS contains the reduced costs $c_N^T - c_B^T B^{-1}N$. In addition to _OBJ_ID_ and _RHS_ID_, it contains the following variables:

- **_BASIC_** gives the names of the basic variables in the solution.
- INVB_R gives the values of $B^{-1}r$, where r is the right-hand-side vector.

STRUCTURAL VARIABLES the values in the tableau, namely $B^{-1}A$.

Input Data Sets

In addition to the DATA= input data set, PROC LP recognizes the ACTIVEIN= and the PRIMALIN= data sets.

ACTIVEIN= Data Set

The ACTIVEIN= data set contains a representation of the current active tree. The format is identical to that of the ACTIVEOUT= data set.

PRIMALIN= Data Set

The format of the PRIMALIN= data set is identical to the PRIMALOUT= data set. PROC LP uses the PRIMALIN= data set to identify variables at their upper bounds in the current solution and variables that are basic in the current solution.

You can add observations to the end of the problem data set if they define cost (righthand-side) sensitivity change vectors and have TYPE=*PRICESEN (RHSSEN)*. This enables you to solve a problem, save the solution in a SAS data set, and perform sensitivity analysis later. You can also use the PRIMALIN= data set to restart problems that have not been completely solved or to which new variables have been added.

Displayed Output

The output from the LP procedure is discussed in the following six sections:

- Problem Summary
- Solution Summary including a Variable Summary and a Constraint Summary
- Infeasible Information Summary
- RHS Sensitivity Analysis Summary (the RHS Range Analysis Summary is not discussed)
- Price Sensitivity Analysis Summary (the Price Range Analysis Summary is not discussed)
- Iteration Log.

For integer-constrained problems, the procedure also displays an Integer Iteration Log. The description of this Log can be found in the "Integer Programming" section on page 99. When you request that the tableau be displayed, the procedure displays the Current Tableau. The description of this can be found in the "The Reduced Costs, Dual Activities, and Current Tableau" section on page 96.

A problem data set can contain a set of constraints with several right-hand sides and several objective functions. PROC LP considers each combination of right-hand side and objective function as defining a new linear programming problem and solves each, performing all specified sensitivity analysis on each problem. For each problem defined, PROC LP displays a new sequence of output sections. Example 3.1 in the "Examples" section on page 123 discusses each of these elements.

The LP procedure produces the following displayed output by default.

The Problem Summary

The problem summary includes the

- type of optimization and the name of the objective row (as identified by the ID or ROW variable)
- name of the SAS variable that contains the right-hand-side constants
- name of the SAS variable that contains the type keywords
- density of the coefficient matrix (the ratio of the number of nonzero elements to the number of total elements) after the slack and surplus variables have been appended
- number of each type of variable in the mathematical program
- number of each type of constraint in the mathematical program

The Solution Summary

The solution summary includes the

- termination status of the procedure
- objective value of the current solution
- number of phase 1 iterations that were completed
- number of phase 2 iterations that were completed
- number of phase 3 iterations that were completed
- number of integer iterations that were completed
- number of integer feasible solutions that were found
- number of initial basic feasible variables identified
- time used in solving the problem excluding reading the data and displaying the solution
- number of inversions of the basis matrix
- current value of several of the options

The Variable Summary

The variable summary includes the

- column number associated with each structural or logical variable in the problem
- name of each structural or logical variable in the problem. (PROC LP gives the logical variables the name of the constraint ID. If no ID variable is specified, the procedure names the logical variable _OBSn_, where *n* is the observation that describes the constraint.)
- variable's status in the current solution. The status can be BASIC, DEGEN, ALTER, blank, LOWBD, or UPPBD, depending upon whether the variable is a basic variable, a degenerate variable, a variable that can be brought into the basis to define an alternate optimal solution, a nonbasic variable at its default lower bound 0, a nonbasic variable at its lower bound, or a nonbasic variable at its upper bound
- kind of variable (whether it is logical or structural, and, if structural, its bound type, or other value restriction)
- value of the objective coefficient associated with each variable
- activity of the variable in the current solution
- variable's reduced cost in the current solution

The Constraint Summary

The constraint summary includes the

- constraint row number and its ID
- the kind of constraint (whether it is an OBJECTIVE, LE, EQ, GE, RANGELE, RANGEEQ, RANGEGE, or FREE row)

- number of the slack or surplus variables associated with the constraint row
- value of the right-hand-side constant associated with the constraint row
- current activity of the row (excluding logical variables)
- current activity of the dual variable (shadow price) associated with the constraint row

The Infeasible Information Summary

The infeasible information summary includes the

- name of the infeasible row or variable
- current activity for the row or the variable
- type of the row or the variable
- value of right-hand-side constant
- name of each nonzero and nonmissing variable in the row
- activity and upper and lower bounds for the variable

The RHS Sensitivity Analysis Summary

The RHS sensitivity analysis summary includes the

- value of ϕ_{min}
- leaving variable when $\phi = \phi_{min}$
- objective value when $\phi = \phi_{min}$
- value of ϕ_{max}
- leaving variable when $\phi = \phi_{max}$
- objective value when $\phi = \phi_{max}$
- column number and name of each logical and structural variable
- variable's status when $\phi \in [\phi_{min}, \phi_{max}]$
- variable's reduced cost when $\phi \in [\phi_{min}, \phi_{max}]$
- value of right-hand-side constant when $\phi = \phi_{min}$
- activity of the variable when $\phi = \phi_{min}$
- value of right-hand-side constant when $\phi = \phi_{max}$
- activity of the variable when $\phi = \phi_{max}$

The Price Sensitivity Analysis Summary

The price sensitivity analysis summary includes the

- value of ϕ_{min}
- entering variable when $\phi = \phi_{min}$
- objective value when $\phi = \phi_{min}$
- value of ϕ_{max}
- entering variable when $\phi = \phi_{max}$

- objective value when $\phi = \phi_{max}$
- column number and name of each logical and structural variable
- variable's status when $\phi \in [\phi_{min}, \phi_{max}]$
- activity of the variable when $\phi \in [\phi_{min}, \phi_{max}]$
- price of the variable when $\phi = \phi_{min}$
- variable's reduced cost when $\phi = \phi_{min}$
- price of the variable when $\phi = \phi_{max}$
- variable's reduced cost when $\phi = \phi_{max}$

The Iteration Log

The iteration log includes the

- phase number
- iteration number in each phase
- name of the leaving variable
- name of the entering variable
- variable's reduced cost
- objective value

ODS Table and Variable Names

PROC LP assigns a name to each table it creates. You can use these names to select output tables when using the Output Delivery System (ODS).

Table Name	Description	Statement/Option
ProblemSummary	Problem summary	default
SolutionSummary	Solution summary	default
VariableSummary	Variable summary	default
ConstraintSummary	Constraint summary	default
IterationLog	Iteration log	FLOW
IntegerIterationLog	Integer iteration log	default
PriceSensitivitySummary	Price sensitivity analysis	default or PRINT (PRICESEN or COL-
	summary	UMN /SENSITIVITY)
PriceActivities	Price activities at min phi and	default or PRINT (PRICESEN or COL-
	max phi	UMN /SENSITIVITY)
PriceActivity	Price activity at min phi or	PRICEPHI= and PARAPRINT
	max phi	
PriceParametricLog	Price parametric program-	PRICEPHI=
	ming log	
PriceRangeSummary	Price range analysis	RANGEPRICE or PRINT RANGEPRICE
RhsSensitivitySummary	RHS sensitivity analysis	default or PRINT (RHSSEN or ROW
	summary	/SENSITIVITY)
RhsActivities	RHS activities at min phi and	default or PRINT (RHSSEN or ROW
	max phi	/SENSITIVITY)
RhsActivity	RHS activity at min phi or	RHSPHI= and PARAPRINT
	max phi	
RhsParametricLog	RHS parametric program-	RHSPHI=
	ming log	
RhsRangeSummary	RHS range analysis	RANGERHS or PRINT RANGERHS
InfeasibilitySummary	Infeasible row or variable	default
	summary	
InfeasibilityActivity	Variables activity in an infea-	default
	sible row	
CurrentTableau	Current tableau	TABLEAUPRINT or PRINT TABLEAU
Matrix	Technological matrix	PRINT MATRIX
MatrixPicture	Technological matrix picture	PRINT MATRIX /PICTURE
MatrixPictureLegend	Technological matrix picture	PRINT MATRIX /PICTURE
	legend	

Table 3.6. ODS Tables Produced in PROC LP

The following table lists the variable names of the preceding tables used in the ODS template of the LP procedure.

Table Name	Variables
VariableSummary	VarName, Status, Type, Price, Activity, ReducedCost
ConstraintSummary	Row, RowName, Type, SSCol, Rhs, Activity, Dual
IterationLog	Phase, Iteration, EnterVar, EnterCol, LeaveVar, LeaveCol, ReducedCost,
	ObjValue
IntegerIterationLog	Iteration, Problem, Condition, Objective, Branch, Value, SumOfInfeas, Active,
	Proximity
PriceActivities	Col, VarName, Status, Activity, MinPrice, MinReducedCost, MaxPrice,
	MaxReducedCost
PriceActivity	Col, VarName, Status, Activity, Price, ReducedCost
PriceParametricLog	LeaveVar, LeaveCol, EnterVar, EnterCol, ObjValue, CurrentPhi
PriceRangeSummary	Col, VarName, MinPrice, MinEnterVar, MinObj, MaxPrice, MaxEnterVar,
	MaxObj
RhsActivities	Col, VarName, Status, ReducedCost, MinRhs, MinActivity, MaxRhs,
	MaxActivity
RhsActivity	Col, VarName, Status, ReducedCost, Rhs, Activity,
RhsParametricLog	LeaveVar, LeaveCol, EnterVar, EnterCol, ObjValue, CurrentPhi
RhsRangeSummary	RowName, MinRhs, MinLeaveVar, MinObj, MaxRhs, MaxLeaveVar, MaxObj
InfeasibilityActivity	VarName, Coefficient, Activity, Lower, Upper

Table 3.7. Variables names for the ODS Tables Produced in PROC LP

Examples

The following fourteen examples illustrate some of the capabilities of PROC LP. These examples, together with the other SAS/OR examples, can be found in the SAS sample library. A description of the features of PROC LP as shown in the examples are

Example 3.1	dense input format
Example 3.2	sparse input format
Example 3.3	the RANGEPRICE option to show you the range over which each objective coefficient can vary without changing the variables in the basis
Example 3.4	more sensitivity analysis and restarting a problem
Example 3.5	parametric programming
Example 3.6	special ordered sets
Example 3.7	goal programming
Example 3.8	integer programming
Example 3.9	an infeasible problem
Example 3.10	restarting integer programs
Example 3.11	controlling the search of the branch and bound tree
Example 3.12	matrix generation and report writing for an assignment problem
Example 3.13	matrix generation and report writing for a scheduling problem
Example 3.14	a multicommodity transshipment problem

Example 3.1. An Oil Blending Problem

The blending problem presented in the introduction is a good example for demonstrating some of the features of the LP procedure. Recall that a step in refining crude oil into finished oil products involves a distillation process that splits crude into various streams. Suppose that there are three types of crude available: Arabian light, Arabian heavy, and Brega. These are distilled into light naphtha, intermediate naphtha, and heating oil. Using one of two recipes, these in turn are blended into jet fuel.

Assume that you can sell as much fuel as is produced. What production strategy maximizes the profit from jet fuel sales? The following SAS code demonstrates a way of answering this question using linear programming. The SAS data set is a representation of the formulation for this model given in the introductory section.

```
data;
   input _row_ $14.
          a_light a_heavy brega naphthal naphthai heatingo jet_1
          jet 2
                 _type_ $ _rhs_;
   datalines;
                   -175 -165 -205
                                         0
profit
                                     0
                                            0
                                                300
                                                     300
                                                          max
                   .035 .030 .045
                                    -1
                                         0
                                            0
                                                 0
                                                      0
                                                                   0
naphtha_l_conv
                                                           eq
naphtha_i_conv
                  .100 .075 .135
                                     0 -1
                                            0
                                                 0
                                                      0
                                                                   0
                                                          eq
                   .390 .300 .430
                                     0
                                         0 -1
heating_o_conv
                                                0
                                                      0
                                                          eq
                                                                   0
                                        .3.7
recipe 1
                      0
                           0
                                 0
                                     0
                                                -1
                                                      0
                                                          eq
                                                                   0
                      0
                           0
                                 0
                                    .2
                                         0.8
                                                 0
recipe_2
                                                     -1
                                                          eq
                                                                   0
                   110
available
                         165
                                80
                                                         upperbd .
                                                .
                                                     .
;
proc lp;
run;
```

The _ROW_ variable contains the names of the rows in the model; the variables A_LIGHT to JET_2 are the names of the structural variables in the model; the _TYPE_ variable contains the keywords that tell the LP procedure how to interpret each row in the model; and the _RHS_ variable gives the value of the right-hand-side constants.

The structural variables are interpreted as the quantity of each type of constituent or finished product. For example, the value of A_HEAVY in the solution is the amount of Arabian heavy crude to buy while the value of JET_1 in the solution is the amount of recipe 1 jet fuel that is produced. As discussed previously, the values given in the model data set are the technological coefficients whose interpretation depends on the model. In this example, the coefficient -175 in the PROFIT row for the variable A_LIGHT gives a cost coefficient (because the row with_ROW_=PROFIT has _TYPE_=MAX) for the structural variable A_LIGHT. This means that for each unit of Arabian heavy crude purchased, a cost of 175 units is incurred.

The coefficients 0.035, 0.100, and 0.390 for the A_LIGHT variable give the percentages of each unit of Arabian light crude that is distilled into the light naphtha, intermediate naphtha, and heating oil components. The 110 value in the row _ROW_=AVAILABLE gives the quantity of Arabian light that is available.

PROC LP produces the following Problem Summary output. Included in the summary is an identification of the objective, defined by the first observation of the problem data set; the right-hand-side variable, defined by the variable _RHS_; and the type identifier, defined by the variable _TYPE_. See Output 3.1.1.
The LP Procedure							
	Problem Summary						
	Objective Function Rhs Variable Type Variable Problem Density (%)	Max profit _rhs_ _type_ 45.00					
	Variables	Number					
	Non-negative Upper Bounded	5 3					
	Total	8					
	Constraints	Number					
	EQ Objective	5 1					
	Total	6					

Output 3.1.1. Problem Summary for the Oil Blending Problem

The next section of output (Output 3.1.2) contains the Solution Summary, which indicates whether or not an optimal solution was found. In this example, the procedure terminates successfully (with an optimal solution), with 1544 as the value of the objective function. Also included in this section of output is the number of phase 1 and phase 2 iterations, the number of variables used in the initial basic feasible solution, and the time used to solve the problem. For several options specified in the PROC LP statement, the current option values are also displayed.

Output 3.1.2. Solution Summary for the Oil Blending Problem

The LP Proce	edure	
Solution Sum	mary	
Terminated Succe	essfully	
Objective Value	1544	
Phase 1 Iterations	0	
Phase 2 Iterations	7	
Phase 3 Iterations	0	
Integer Iterations	0	
Integer Solutions	0	
Initial Basic Feasible Varia	ables 2	
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	

126 • Chapter 3. The LP Procedure

The next section of output (Output 3.1.3) contains the Variable Summary. A line is displayed for each variable in the mathematical program with the variable name, the status of the variable in the solution, the type of variable, the variable's price coefficient, the activity of the variable in the solution, and the reduced cost for the variable. The status of a variable can be

BASIC	if the variable is a basic variable in the solution
DEGEN	if the variable is a basic variable whose activity is at its input lower bound
ALTER	if the variable is nonbasic and is basic in an alternate optimal solution
LOWBD	if the variable is nonbasic and is at its lower bound
UPPBD	if the variable is nonbasic and is at its upper bound

The TYPE column shows how PROC LP interprets the variable in the problem data set. Types include the following:

NON-NEG	if the variable is a nonnegative variable with lower bound 0 and upper bound <i>+INFINITY</i>
LOWERBD	if the variable has a lower bound specified in a <i>LOWERBD</i> observation and upper bound <i>+INFINITY</i>
UPPERBD	if the variable has an upper bound that is less than <i>+INFINITY</i> and lower bound 0 This upper bound is specified in an <i>UPPERBD</i> observation
UPLOWBD	if the variable has a lower bound specified in a <i>LOWERBD</i> observation and an upper bound specified in an <i>UPPERBD</i> observation
INTEGER	if the variable is constrained to take integer values. If this is the case, then it must also be upper and lower bounded
BINARY	if the variable is constrained to take value 0 or 1
UNRSTRCT	if the variable is an unrestricted variable having bounds of - <i>INFINITY</i> and <i>+INFINITY</i>
SLACK	if the variable is a slack variable that PROC LP has appended to a <i>LE</i> constraint. For variables of this type, the variable name is the same as the name of the constraint (given in the ROW variable) for which this variable is the slack. A nonzero slack variable indicates that the constraint is not tight. The slack is the amount by which the right-hand side of the constraint exceeds the left-hand side.
SURPLUS	if the variable is a surplus variable that PROC LP has appended to a <i>GE</i> constraint. For variables of this type, the variable name is the same as the name of the constraint (given in the ROW variable) for which this variable is the surplus. A nonzero surplus variable indicates that the constraint is not tight. The surplus is the amount by which the left-hand side of the constraint exceeds the right-hand side.

The Variable Summary gives the value of the structural variables at optimality. In this example, it tells you how to produce the jet fuel to maximize your profit. You should buy 110 units of A_LIGHT and 80 units of BREGA. These are used to make 7.45 units of NAPHTHAL, 21.8 units of NAPHTHAI, and 77.3 units of HEATINGO. These in turn are used to make 60.65 units of JET_1 using recipe 1 and 63.33 units of JET_2 using recipe 2.

Output 3.1.3. Variable Summary for the Oil Blending Problem

The LP Procedure								
Variable Summary								
	Reduced							
Col	Variable Name	Status	Туре	Price	Activity	Cost		
1	a_light	UPPBD	UPPERBD	-175	110	11.6		
2	a_heavy		UPPERBD	-165	0	-21.45		
3	brega	UPPBD	UPPERBD	-205	80	3.35		
4	naphthal	BASIC	NON-NEG	0	7.45	0		
5	naphthai	BASIC	NON-NEG	0	21.8	0		
6	heatingo	BASIC	NON-NEG	0	77.3	0		
7	jet_1	BASIC	NON-NEG	300	60.65	0		
8	jet_2	BASIC	NON-NEG	300	63.33	0		

The reduced cost associated with each nonbasic variable is the marginal value of that variable if it is brought into the basis. In other words, the objective function value would (assuming no constraints were violated) increase by the reduced cost of a nonbasic variable if that variable's value increased by one. Similarly, the objective function value would (assuming no constraints were violated) decrease by the reduced cost of a nonbasic variable if that variable's value were decreased by one. Basic variables always have a zero reduced cost. At optimality, for a maximization problem, nonbasic variables that are not at an upper bound have nonpositive reduced costs (for example, A_HEAVY has a reduced cost of -21.45). The objective would decrease if they were to increase beyond their optimal values. Nonbasic variables at upper bounds have nonnegative reduced costs, showing that increasing the upper bound (if the reduced cost is not zero) does not decrease the objective. For a nonbasic variable at its upper bound, the reduced cost is the marginal value of increasing its upper bound, often called its shadow price.

For minimization problems, the definition of reduced costs remain the same but the conditions for optimality change. For example, at optimality the reduced costs of all non upper-bounded variables are nonnegative, and the reduced costs of upper-bounded variables at their upper bound are nonpositive.

The next section of output (Output 3.1.4) contains the Constraint Summary. For each constraint row, free row, and objective row, a line is displayed in the Constraint Summary. Included on the line are the constraint name, the row type, the slack or surplus variable associated with the row, the right-hand-side constant associated with the row, the activity of the row (not including the activity of the slack and surplus variables), and the dual activity (shadow prices).

A dual variable is associated with each constraint row. At optimality, the value of this variable, the dual activity, tells you the marginal value of the right-hand-side constant. For each unit increase in the right-hand-side constant, the objective changes by this amount. This quantity is also known as the shadow price. For example, the marginal value for the right-hand-side constant of constraint HEATING_O_CONV is -450.0.

Output 3.1.4. Constraint Summary for the Oil Blending Problem

	The L	P Procedur	e				
Constraint Summary							
Constra	aint	S/S			Dual		
Row Name	Туре	Col	Rhs	Activity	Activity		
1 profit	OBJECTVE		0	1544			
2 naphtha	a_l_conv EQ		0	0	-60		
3 naphth	a_i_conv EQ		0	0	-90		
4 heating	g_o_conv EQ		0	0	-450		
5 recipe	_1 EQ		0	0	-300		
6 recipe	_2 EQ	•	0	0	-300		

Example 3.2. A Sparse View of the Oil Blending Problem

Typically, mathematical programming models are very sparse. This means that only a small percentage of the coefficients are nonzero. The sparse problem input is ideal for these models. The oil blending problem in the "An Introductory Example" section has a sparse form. This example shows the same problem in a sparse form with the data given in a different order. In addition to representing the problem in a concise form, the sparse format

- allows long column names
- enables easy matrix generation (See Example 3.12, Example 3.13, and Example 3.14.)
- is compatible with MPS sparse format

The model in the sparse format is solved by invoking PROC LP with the SPARSE-DATA option as shown below.

```
data oil;
   input _type_ $8. _col_ $14. _row_ $17. _coef_ ;
   datalines;
                       profit
max
        arabian_light profit
                                         -175
        arabian_heavy profit
                                         -165
        brega
                      profit
                                         -205
        jet 1
                      profit
                                          300
.
        jet_2
                      profit
                                          300
                      napha_l_conv
eα
        .
        arabian_light napha_l_conv
                                         .035
```

•	arabian_heavy	napha_l_conv	.030
•	brega	napha_l_conv	.045
•	naphtha_light	napha_l_conv	-1
eq	•	napha_i_conv	•
•	arabian_light	napha_i_conv	.100
•	arabian_heavy	napha_i_conv	.075
•	brega	napha_i_conv	.135
•	naphtha_inter	napha_i_conv	-1
eq	•	heating_oil_conv	•
•	arabian_light	heating_oil_conv	.390
•	arabian_heavy	heating_oil_conv	.300
•	brega	heating_oil_conv	.430
•	heating_oil	heating_oil_conv	-1
eq	•	recipe_1	•
•	naphtha_inter	recipe_1	.3
•	heating_oil	recipe_1	.7
eq	•	recipe_2	•
•	jet_1	recipe_1	-1
•	naphtha_light	recipe_2	.2
•	heating_oil	recipe_2	.8
•	jet_2	recipe_2	-1
upperbd	•	available	•
•	arabian_light	available	110
•	arabian_heavy	available	165
•	brega	available	80
;			
proc lp run;	SPARSEDATA;		

The output from PROC LP follows.

Output 3.2.1. Output for the Sparse Oil Blending Problem

The LP Procedure							
	Problem Summary						
	Objective Function Rhs Variable Type Variable	Max profit _rhs_ _type_					
	Problem Density (%) Variables	45.00 Number					
	Non-negative Upper Bounded	5 3					
	Total	8					
	Constraints	Number					
	EQ Objective	5 1					
	Total	6					

The LP Procedure						
Solution Summary						
Terminated Successfully	7					
Objective Value	1544					
Phase 1 Iterations	0					
Phase 2 Iterations	7					
Phase 3 Iterations	0					
Integer Iterations	0					
Integer Solutions	0					
Initial Basic Feasible Variables	2					
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					

The LP Procedure							
Variable Summary							
Col Variable	Name Status	Туре	Price	Activity	Reduced Cost		
1 arabian_h	neavy	UPPERBD	-165	0	-21.45		
2 arabian_1	ight UPPBD	UPPERBD	-175	110	11.6		
3 brega	UPPBD	UPPERBD	-205	80	3.35		
4 heating_c	il BASIC	NON-NEG	0	77.3	0		
5 jet_1	BASIC	NON-NEG	300	60.65	0		
6 jet_2	BASIC	NON-NEG	300	63.33	0		
7 naphtha_i	nter BASIC	NON-NEG	0	21.8	0		
8 naphtha_1	ight BASIC	NON-NEG	0	7.45	0		

Objective Coefficients • 13

The LP Procedure								
Constraint Summary								
Row	Constraint Name	Type	s/s Col	Rhs	Activity	Dual		
1.01			001		1544	1100111207		
T	profit	OBJECIVE	•	0	1544	•		
2	napha_l_conv	EQ	•	0	0	-60		
3	napha_i_conv	EQ	•	0	0	-90		
4	heating_oil_conv	EQ	•	0	0	-450		
5	recipe_1	EQ	•	0	0	-300		
6	recipe_2	EQ	•	0	0	-300		

Example 3.3. Analyzing the Sensitivity of the Solution to Changes in the Objective Coefficients

Simple solution of a linear program is often not enough. A manager needs to evaluate how sensitive the solution is to changing assumptions. The LP procedure provides several tools that are useful for "what if," or sensitivity, analysis. One tool studies the effects of changes in the objective coefficients.

For example, in the oil blending problem, the cost of crude and the selling price of jet fuel can be highly variable. If you want to know the range over which each objective coefficient can vary without changing the variables in the basis, you can use the RANGEPRICE option in the PROC LP statement.

proc lp sparsedata RANGEPRICE PRIMALOUT=SOLUTION; run;

In addition to the Problem and Solution summaries, the LP procedure produces a Price Range Summary, shown in Output 3.3.1.

For each structural variable, the upper and lower ranges of the price (objective function coefficient) and the objective value are shown. The blocking variables, those variables that would enter the basis if the objective coefficient were perturbed further, are also given. For example, the output shows that if the cost of ARABIAN_LIGHT crude were to increase from 175 to 186.6 per unit (remember that you are maximizing profit so the ARABIAN_LIGHT objective coefficient would decrease from -175 to -186.6), then it would become optimal to use less of this crude for any fractional increase in its cost. Increasing the unit cost to 186.6 would drive its reduced cost to zero. Any additional increase would drive its reduced cost negative and would destroy the optimality conditions; thus, you would want to use less of it in your processing. The output shows that, at the point where the reduced cost is zero, you would only be realizing a profit of 268 = 1544 - (110 * 11.6) and that ARABIAN_LIGHT enters the basis, that is, leaves its upper bound. On the other hand, if the cost of ARABIAN_HEAVY were to decrease to 143.55, you would want to stop using the formulation of 110 units of ARABIAN_LIGHT and 80 units of BREGA and switch to a production scheme that included ARABIAN_HEAVY, in which case the profit would increase from the 1544 level.

131

		The LP Proce	edure						
	Price Range Analysis								
			Minimum Phi-						
Col	Variable Name	Price	Entering	Objective					
1	arabian_heavy	-INFINITY		1544					
2	arabian_light	-186.6	arabian_light	268					
3	brega	-208.35	brega	1276					
4	heating_oil	-7.790698	brega	941.77907					
5	jet_1	290.19034	brega	949.04392					
6	jet_2	290.50992	brega	942.99292					
7	naphtha_inter	-24.81481	brega	1003.037					
8	naphtha_light	-74.44444	brega	989.38889					
	Pi	rice Range Ar	nalysis						
		Maximum H	?hi						
	Price	Entering	Objective						
	-143.55	arabian_heav	<i>r</i> y 1544						
	INFINITY	. –	INFINITY						
	INFINITY		INFINITY						
	71.5	arabian heav	<i>r</i> y 7070.95						
	392.25806	arabian heav	ry 7139.4516						
	387.19512	arabian heav	ry 7066.0671						
	286	arabian heav	ry 7778.8						
	715	arabian heav	ry 6870.75						
		_	-						

Output 3.3.1. Price Range Summary for the Oil Blending Problem

Note that in the PROC LP statement, the PRIMALOUT=SOLUTION option was given. This caused the procedure to save the optimal solution in a SAS data set named SOLUTION. This data set can be used to perform further analysis on the problem without having to restart the solution process. Example 3.4 shows how this is done. A display of the data follows in Output 3.3.2.

Output 3.3.2. The PRIMALOUT= Data Set for the Oil Blending Problem

	0	R			s	L	_		_	R
	в	н		_	т	в	v	в	P	_
	J	S	_	т	A	о	A	0	R	C
	_	_	v	Y	т	U	L	U	I	0
0	I	I	A	P	υ	N	υ	N	C	S
b	D	D	R	E	S	D	E	D	Е	т
s	_	_	_	_	_	_	_	_	_	_
1	profit	$_{rhs}$	arabian_heavy	UPPERBD		0	0.00	165	-165	-21.45
2	profit	_rhs_	arabian_light	UPPERBD	_UPPER_	0	110.00	110	-175	11.60
3	profit	_rhs_	brega	UPPERBD	_UPPER_	0	80.00	80	-205	3.35
4	profit	_rhs_	heating_oil	NON-NEG	_BASIC_	0	77.30	1.7977E308	0	0.00
5	profit	_rhs_	jet_1	NON-NEG	_BASIC_	0	60.65	1.7977E308	300	0.00
6	profit	_rhs_	jet_2	NON-NEG	_BASIC_	0	63.33	1.7977E308	300	0.00
7	profit	_rhs_	naphtha_inter	NON-NEG	_BASIC_	0	21.80	1.7977E308	0	0.00
8	profit	_rhs_	naphtha_light	NON-NEG	_BASIC_	0	7.45	1.7977E308	0	0.00
9	profit	_rhs_	PHASE_1_OBJECTIV	OBJECT	_DEGEN_	0	0.00	0	0	0.00
10	profit	_rhs_	profit	OBJECT	_BASIC_	0	1544.00	1.7977E308	0	0.00

Example 3.4. Additional Analysis of the Sensitivity of the Solution to Changes in the Objective Coefficients

The objective coefficient ranging analysis, discussed in the last example, is useful for accessing the effects of changing costs and returns on the optimal solution if each objective function coefficient is modified in isolation. However, this is often not the case.

Suppose you anticipate that the cost of crude will be increasing and you want to examine how that will affect your optimal production plans. Furthermore, you estimate that if the price of ARABIAN_LIGHT goes up by 1 unit, then the price of ARABIAN_HEAVY will rise by 1.2 units and the price of BREGA will increase by 1.5 units. However, you plan on passing some of your increased overhead on to your jet fuel customers, and you decide to increase the price of jet fuel 1 unit for each unit of increased cost of ARABIAN_LIGHT.

An examination of the solution sensitivity to changes in the cost of crude is a twostep process. First, add the information on the proportional rates of change in the crude costs and the jet fuel price to the problem data set. Then, invoke the LP procedure. The following program accomplishes this. First, it adds a new row, named CHANGE, to the model. It gives this row a type of *PRICESEN*. That tells PROC LP to perform objective function coefficient sensitivity analysis using the given rates of change. The program then invokes PROC LP to perform the analysis. Notice that the PRIMALIN=SOLUTION option is used in the PROC LP statement. This tells the LP procedure to use the saved solution. Although it is not necessary to do this, it will eliminate the need for PROC LP to resolve the problem and can save computing time.

```
data sen;
   input _type_ $ @10 _col_ $13. @24 _row_ $ _coef_;
   datalines;
pricesen .
                        change
         arabian light change
                                  1
         arabian_heavy change 1.2
•
         brega
                        change 1.5
•
         jet_1
                        change
                                -1
•
         jet 2
                        change
                                -1
.
;
data;
   set oil sen;
proc lp sparsedata primalin=solution;
run;
```

Output 3.4.1 shows the range over which the current basic solution remains optimal so that the current production plan need not change. The objective coefficients are modified by adding ϕ times the change vector given in the SEN data set, where ϕ ranges from a minimum of -4.15891 to a maximum of 29.72973. At the minimum value of ϕ , the profit decreases to 1103.073. This value of ϕ corresponds to an increase in the cost of ARABIAN_HEAVY to 169.99 (namely, $-175 + \phi * 1.2$), ARABIAN_LIGHT

133

134 • Chapter 3. The LP Procedure

to 179.158 (= $-175 + \phi * 1$), and BREGA to 211.23 (= $-205 + \phi * 1.5$), and corresponds to an increase in the price of JET_1 and JET_2 to 304.15 (= $300 + \phi * -1$). These values can be found in the Price column under the section labeled Minimum Phi.

Output 3.4.1. The Price Sensitivity Analysis Summary for the Oil Blending Problem

	The LP Procedure									
		Summary ange								
		Minimum H	?hi	-4.	.158907511					
		Entering	Variable		brega					
		Optimal (Dbjective	110)3.0726257					
		Mawimum I	Dhi	20	72072073					
		Entering	Variable	aral	nian heavy					
		Ontimal (biective	469	35 9459459					
		opermar (Djective	102	,5.,715,155					
				Minim	um Phi	Maximu	um Phi			
					Reduced		Reduced			
Col	Variable Name	Status	Activity	Price	Cost	Price	Cost			
1	arabian_heavy		0	-169.9907	-24.45065	-129.3243	0			
2	arabian_light	UPPBD	110	-179.1589	10.027933	-145.2703	22.837838			
3	brega	UPPBD	80	-211.2384	0	-160.4054	27.297297			
4	heating_oil	BASIC	77.3	0	0	0	0			
5	jet_1	BASIC	60.65	304.15891	0	270.27027	0			
6	jet_2	BASIC	63.33	304.15891	0	270.27027	0			
7	naphtha_inter	BASIC	21.8	0	0	0	0			
8	naphtha_light	BASIC	7.45	0	0	0	0			

The Price Sensitivity Analysis Summary also shows the effects of lowering the cost of crude and lowering the price of jet fuel. In particular, at the maximum ϕ of 29.72973, the current optimal production plan yields a profit of 4695.946. Any increase or decrease in ϕ beyond the limits given results in a change in the production plan. More precisely, the columns that constitute the basis change.

Example 3.5. Price Parametric Programming for the Oil Blending Problem

This example continues to examine the effects of a change in the cost of crude and the selling price of jet fuel. Suppose that you know the cost of ARABIAN_LIGHT crude is likely to increase 30 units, with the effects on oil and fuel prices as described in Example 3.4. The analysis in the last example only accounted for an increase of a little over 4 units (because the minimum ϕ was -4.15891). Because an increase in the cost of ARABIAN_LIGHT beyond 4.15891 units requires a change in the optimal basis, it may require a change in the optimal production strategy as well. This type of analysis, where you want to find how the solution changes with changes in the objective function coefficients or right-hand-side vector, is called parametric programming. You can answer this question by using the PRICEPHI= option in the PROC LP statement. The following program instructs PROC LP to continually increase the cost of the crudes and the return from jet fuel using the ratios given previously, until the cost of ARABIAN_LIGHT increases at least 30 units.

proc lp sparsedata primalin=solution PRICEPHI=-30; run;

The PRICEPHI= option in the PROC LP statement tells PROC LP to perform parametric programming on any price change vectors specified in the problem data set. The value of the PRICEPHI= option tells PROC LP how far to change the value of ϕ and in what direction. A specification of PRICEPHI=-30 tells PROC LP to continue pivoting until the problem has objective function equal to the original objective function value -30 * change vector.

Output 3.5.1 shows the result of this analysis. The first page is the Price Sensitivity Analysis Summary, as discussed in Example 3.4. The next page is an accounting for the change in basis as a result of decreasing ϕ beyond -4.1589. It shows that BREGA left the basis at an upper bound and entered the basis at a lower bound. The interpretation of these basis changes can be difficult (Hadley 1962; Dantzig 1963).

The last page of output shows the optimal solution at the displayed value of ϕ , namely -30.6878. At an increase of 30.6878 units in the cost of ARABIAN_LIGHT and the related changes to the other crudes and the jet fuel, it is optimal to modify the production of jet fuel as shown in the activity column. Although this plan is optimal, it results in a profit of 0. This may suggest that the ratio of a unit increase in the price of jet fuel for unit increase in the cost of ARABIAN_LIGHT is lower than desirable.

Output 3.5.1. Price Parametric Programming for the Oil Blending Problem

	The LP Procedure									
		Minimum 1	Phi	-4	158907511					
		Entering	Variable		brega					
		Optimal (Objective	110	03.0726257					
		Maximum 1	Phi	29	9.72972973					
		Entering	Variable	aral	oian_heavy					
		Optimal (Objective	469	95.9459459					
				Minimu	um Phi	Maximu	um Phi			
					Reduced		Reduced			
Col	Variable Name	Status	Activity	Price	Cost	Price	Cost			
1	arabian_heavy		0	-169.9907	-24.45065	-129.3243	0			
2	arabian_light	UPPBD	110	-179.1589	10.027933	-145.2703	22.837838			
3	brega	UPPBD	80	-211.2384	0	-160.4054	27.297297			
4	heating_oil	BASIC	77.3	0	0	0	0			
5	jet_1	BASIC	60.65	304.15891	0	270.27027	0			
6	jet_2	BASIC	63.33	304.15891	0	270.27027	0			
7	naphtha_inter	BASIC	21.8	0	0	0	0			
8	naphtha_light	BASIC	7.45	0	0	0	0			

	The LP Procedure	9		
Price Se	Parametric Program nsitivity Vector (mming Log change		
Leaving Variable	Entering Variable	Objective	Current Phi	
brega	brega	1103.0726 -	-4.158908	

The LP Procedure										
Price Sensitivity Analysis Summary										
	Sensi	tivity \	Vector cha	ange						
	Minimum Phi		-30	.68783069						
	Entering Va	riable	aral	oian_light						
	Optimal Obje	ective		0						
					-1.1					
				Minimu	um Phi					
					Reduced					
Col	Variable Name	Status	Activity	Price	Cost					
1	arabian_heavy		0	-201.8254	-43.59127					
2	arabian_light	ALTER	110	-205.6878	0					
3	brega		0	-251.0317	-21.36905					
4	heating_oil	BASIC	42.9	0	0					
5	jet_1	BASIC	33.33	330.68783	0					
6	jet_2	BASIC	35.09	330.68783	0					
7	naphtha_inter	BASIC	11	0	0					
8	naphtha_light	BASIC	3.85	0	0					

What is the optimal return if ϕ is exactly -30? Because the change in the objective is linear as a function of ϕ , you can calculate the objective for any value of ϕ between those given by linear interpolation. For example, for any ϕ between -4.1589 and -30.6878, the optimal objective value is

 $\phi * (1103.0726 - 0) / (-4.589 - 30.6878) + b$

where

b = 30.6878 * (1103.0726 - 0)/(-4.589 - 30.6878)

For $\phi = -30$, this is 28.5988.

Example 3.6. Special Ordered Sets and the Oil Blending Problem

Often managers want to evaluate the cost of making a choice among alternatives. In particular, they want to make the most profitable choice. Suppose that only one oil crude can be used in the production process. This identifies a set of variables of which only one can be above its lower bound. This additional restriction could be included in the model by adding a binary integer variable for each of the three crudes. Constraints would be needed that would drive the appropriate binary variable to 1 whenever the corresponding crude is used in the production process. Then a constraint limiting the total of these variables to only one would be added. A similar formulation for a fixed charge problem is shown in Example 3.8.

The *SOSLE* type implicitly does this. The following DATA step adds a row to the model that identifies which variables are in the set. The *SOSLE* type tells the LP procedure that only one of the variables in this set can be above its lower bound. If you use the *SOSEQ* type, it tells PROC LP that exactly one of the variables in the set must be above its lower bound. Only integer variables can be in an *SOSEQ* set.

```
data special;
   input _type_ $ @7 _col_ $13. @21 _row_ $ _coef_;
   datalines;
SOSLE .
                    special .
      arabian_light special 1
.
      arabian_heavy special 1
                   special 1
      brega
•
;
data;
   set oil special;
proc lp sparsedata;
run;
```

Output 3.6.1 includes an Integer Iteration Log. This log shows the progress that PROC LP is making in solving the problem. This is discussed in some detail in Example 3.8.

The LP Procedure							
Problem Summary							
	Objective Function	Max profit					
	Rhs Variable	_rhs_					
	Type Variable	_type_					
	Problem Density (%)	45.00					
	Variables	Number					
	Non-negative	5					
	Upper Bounded	3					
	Total	8					
	Constraints	Number					
	EQ	5					
	Objective	1					
	Total	6					

Output 3.6.1. The Oil Blending Problem with a Special Ordered Set

			Th€	≥ LP Procedure				
			Integ	ger Iteration Log	3			
Iter	Problem	Condition	Objective	Branched	Value	Sinfeas	Active	Proximity
1	0	ACTIVE	1544	arabian_light	110	0	2	•
2	-1	SUBOPTIMAL	1276	•		•	1	268
3	1	FATHOMED	268	•	•	•	0	

The LP Procedure		
Solution Summary		
Integer Optimal Soluti	on	
Objective Value	1276	
Phase 1 Iterations	0	
Phase 2 Iterations	7	
Phase 3 Iterations	0	
Integer Iterations	3	
Integer Solutions	1	
Initial Basic Feasible Variables	2	
Time Used (seconds)	0	
Number of Inversions	4	
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	

The LP Procedure										
	Variable Summary									
Col	Variable Name	Status	Type	Price	Activity	Reduced				
001	Variable Mame	BLACUS	TYPE	FIICE	ACCIVICI	COBE				
1	arabian_heavy		UPPERBD	-165	0	-21.45				
2	arabian_light	UPPBD	UPPERBD	-175	110	11.6				
3	brega		UPPERBD	-205	0	3.35				
4	heating_oil	BASIC	NON-NEG	0	42.9	0				
5	jet_1	BASIC	NON-NEG	300	33.33	0				
6	jet_2	BASIC	NON-NEG	300	35.09	0				
7	naphtha_inter	BASIC	NON-NEG	0	11	0				
8	naphtha_light	BASIC	NON-NEG	0	3.85	0				

The LP Procedure									
Constraint Summary									
			s/s			Dual			
Row	Constraint Name	Туре	Col	Rhs	Activity	Activity			
1	profit	OBJECTVE	•	0	1276	•			
2	napha_1_conv	EQ	•	0	0	-60			
3	napha_i_conv	EQ	•	0	0	-90			
4	heating_oil_conv	EQ	•	0	0	-450			
5	recipe_1	EQ	•	0	0	-300			
6	recipe_2	EQ	•	0	0	-300			

The solution shows that only the ARABIAN_LIGHT crude is purchased. The requirement that only one crude be used in the production is met, and the profit is 1276. This tells you that the value of purchasing crude from an additional source, namely BREGA, is worth 1544 - 1276 = 268.

Example 3.7. Goal-Programming a Product Mix Problem

This example shows how to use PROC LP to solve a linear goal-programming problem. PROC LP has the ability to solve a series of linear programs, each with a new objective function. These objective functions are ordered by priority. The first step is to solve a linear program with the highest priority objective function constrained only by the formal constraints in the model. Then, the problem with the next highest priority objective function is solved, constrained by the formal constraints in the model and by the value that the highest priority objective function realized. That is, the second problem optimizes the second highest priority objective function among the alternate optimal solutions to the first optimization problem. The process continues until a linear program is solved for each of the objectives.

This technique is useful for differentiating among alternate optimal solutions to a linear program. It also fits into the formal paradigm presented in goal programming. In goal programming, the objective functions typically take on the role of driving a linear function of the structural variables to meet a target level as closely as possible. The details of this can be found in many books on the subject, including Ignizio (1976).

Consider the following problem taken from Ignizio (1976). A small paint company manufactures two types of paint, latex and enamel. In production, the company uses 10 hours of labor to produce 100 gallons of latex and 15 hours of labor to produce 100 gallons of enamel. Without hiring outside help or requiring overtime, the company has 40 hours of labor available each week. Furthermore, each paint generates a profit at the rate of \$1.00 per gallon. The company has the following objectives listed in decreasing priority.

- avoid the use of overtime
- achieve a weekly profit of \$1000
- produce at least 700 gallons of enamel paint each week

The program to solve this problem follows:

```
data object;
   input _row_ $ latex enamel n1 n2 n3 p1 p2 p3 _type_ $ _rhs_;
   datalines;
overtime
                             1
                                    . min
                                             1
profit
                       1
                                             2
                              .
                                    .
                                      min
                          1
                                             3
enamel
                                      min
           .
                .
                       .
                   .
                             .
overtime 10
               15
                   1
                            -1
                                            40
                          •
                                      eq
                       .
                                      eq 1000
profit
         100 100
                   •
                       1
                          .
                               -1
enamel
                1
                       •
                          1
                                 . -1 eq
                                             7
;
proc lp GOALPROGRAM;
run;
```

The data set called OBJECT contains the model. Its first three observations are the objective rows, and the next three observations are the constraints. The values in the right-hand-side variable _RHS_ in the objective rows give the priority of the objectives. The objective in the first observation with _ROW_=OVERTIME has the highest priority, the objective named PROFIT has the next highest, and the objective named ENAMEL has the lowest. Note that the value of the right-hand-side variable determines the priority, not the order, in the data set.

Because this example is set in the formal goal-programming scheme, the model has structural variables representing negative (N1, N2, and N3) and positive (P1, P2, and P3) deviations from target levels. For example, N1+P1 is the deviation from the objective of avoiding the use of overtime and underusing the normal work time, namely using exactly 40 work hours. The other objectives are handled similarly.

Notice that the PROC LP statement includes the GOALPROGRAM option. Without this option, the procedure would solve three separate problems: one for each of the three objective functions. In that case, however, the procedure would not constrain the second and third programs using the results of the preceding programs; also, the values 1, 2, and 3 for _RHS_ in the objective rows would have no effect.

Output 3.7.1 shows the solution of the goal program, apparently as three linear program outputs. However, examination of the constraint summaries in the second and third problems show that the constraints labeled by the objectives OVERTIME and PROFIT have type *FIXEDOBJ*. This indicates that these objective rows have become constraints in the subsequent problems.

Output 3.7.1. Goal Programming

The LP Proce	edure
Problem Sur	nmary
Objective Function	Min overtime
Rhs Variable	_rhs_
Type Variable	_type_
Problem Density (%)	45.83
Variables	Number
Non-negative	8
Total	8
Constraints	Number
EQ	3
Objective	3
Total	6

The LP Procedure							
Solution Summary							
Terminated Successfully							
Objective Value	0						
Phase 1 Iterations	2						
Phase 2 Iterations	0						
Phase 3 Iterations	0						
Integer Iterations	0						
Integer Solutions	0						
Initial Basic Feasible Variables	5						
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						

	2	he LP Pr	ocedure			
	V	Variable	Summary			
Va	ariable				Reduced	
Col Na	ame Status	Туре	Price	Activity	Cost	
1 la	atex ALTER	NON-NEG	0	0	0	
2 er	namel ALTER	NON-NEG	0	0	0	
3 n1	1 BASIC	NON-NEG	0	40	0	
4 n2	2 BASIC	NON-NEG	0	1000	0	
5 n3	3 BASIC	NON-NEG	0	7	0	
6 p1	1	NON-NEG	1	0	1	
7 p2	2 ALTER	NON-NEG	0	0	0	
8 p3	3 ALTER	NON-NEG	0	0	0	

		The I	LP Proce	dure		
Constraint Summary						
	Constraint		s/s			Dual
Row	Name	Туре	Col	Rhs	Activity	Activity
1	overtime	OBJECTVE		0	0	
2	profit	FREE_OBJ	•	0	1000	
3	enamel	FREE_OBJ	•	0	7	
4	overtime	EQ	•	40	40	0
5	profit	EQ	•	1000	1000	0
6	enamel	EO		7	7	0

The LP Proced	lure	
Problem Summ	ary	
Objective Function Rhs Variable	Min profit _rhs_	
Type variable	_cype_	
Problem Density (%)	45.83	
Variables	Number	
Non-negative	8	
Total	8	
Constraints	Number	
EQ	3	
Objective	3	
Total	6	

The LP Procedure		
Solution Summary		
Terminated Successfully	7	
Objective Value	600	
Phase 1 Iterations	0	
Phase 2 Iterations	2	
Phase 3 Iterations	0	
Integer Iterations	0	
Integer Solutions	0	
Initial Basic Feasible Variables	5	
Time Used (seconds)	0	
Number of Inversions	4	
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	

			The LP Pro	ocedure		
			Variable :	Summary		
	Variable					Reduced
Col	Name	Status	Туре	Price	Activity	Cost
1	latex	BASIC	NON-NEG	0	4	0
2	enamel		NON-NEG	0	0	50
3	nl		NON-NEG	0	0	10
4	n2	BASIC	NON-NEG	1	600	0
5	n3	BASIC	NON-NEG	0	7	0
6	p1	DEGEN	NON-NEG	0	0	0
7	p2		NON-NEG	0	0	1
8	р3	ALTER	NON-NEG	0	0	0

	The LP Procedure						
		Constr	aint Sum	mary			
Row	Constraint Name	Туре	s/s Col	Rhs	Activity	Dual Activity	
1	overtime	FIXEDOBJ	•	0	0	•	
2	profit	OBJECTVE	•	0	600		
3	enamel	FREE_OBJ	•	0	7		
4	overtime	EQ	•	40	40	-10	
5	profit	EQ		1000	1000	1	
6	enamel	EQ	•	7	7	0	

The LP Procedu	re	
Problem Summa	ry	
Objective Function Rhs Variable Type Variable Problem Density (%)	Min enamel _rhs_ _type_ 45.83	
Variables	Number	
Non-negative	8	
Total	8	
Constraints	Number	
EQ Objective	3 3	
Total	6	

The LP Procedure		
Solution Summary		
Terminated Successfully		
Objective Value	7	
Phase 1 Iterations	0	
Phase 2 Iterations	1	
Phase 3 Iterations	0	
Integer Iterations	0	
Integer Solutions	0	
Initial Basic Feasible Variables	5	
Time Used (seconds)	0	
Number of Inversions	6	
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	

		:	The LP Proce	edure			
			Variable Sur	mary			
	Variable					Reduced	
Col	Name	Status	Туре	Price	Activity	Cost	
1	latex	BASIC	NON-NEG	0	4	0	
2	enamel	DEGEN	NON-NEG	0	0	0	
3	nl		NON-NEG	0	0	0.2	
4	n2	BASIC	NON-NEG	0	600	0	
5	n3	BASIC	NON-NEG	1	7	0	
6	pl	DEGEN	NON-NEG	0	0	0	
7	p2		NON-NEG	0	0	0.02	
8	p3		NON-NEG	0	0	1	

		The LI	Procedu:	re			
		Constra	aint Summa	ary			
Row	Constraint Name	Туре	s/s Col	Rhs	Activity	Dual Activity	
1	overtime	FIXEDOBJ	•	0	0		
3	enamel overtime	OBJECTVE	•	0 40	7 40	-0.2	
5	profit enamel	EQ EQ	•	1000 7	1000 7	0.02	
		~					

The solution to the last linear program shows a value of 4 for the variable LATEX and a value of 0 for the variable ENAMEL. This tells you that the solution to the linear goal program is to produce 400 gallons of latex and no enamel paint.

The values of the objective functions in the three linear programs tell you whether you can achieve the three objectives. The activities of the constraints labeled OVER-TIME, PROFIT, and ENAMEL tell you values of the three linear program objectives. Because the first linear programming objective OVERTIME is 0, the highest priority objective, which is to avoid using additional labor, is accomplished. However, because the second and third objectives are nonzero, the second and third priority objectives are not satisfied completely. The PROFIT objective is 600. Because the PROFIT objective is to minimize the negative deviation from the profit constraint, this means that a profit of only 400 = 1000-600 is realized. Similarly, the ENAMEL objective is 7, indicating that there is a negative deviation from the ENAMEL target of 7 units.

Example 3.8. A Simple Integer Program

Recall the linear programming problem presented in the "Introduction to Mathematical Programming" chapter. In that problem, a firm produces two products, chocolates and gumdrops, that are processed by four processes: cooking, color/flavor, condiments, and packaging. The objective is to determine the product mix that maximizes the profit to the firm while not exceeding manufacturing capacities. The problem is extended to demonstrate a use of integer-constrained variables.

Suppose that you must manufacture only one of the two products. In addition, there is a setup cost of 100 if you make the chocolates and 75 if you make the gumdrops. To identify which product will maximize profit, you define two zeroone integer variables, ICHOCO and IGUMDR, and you also define two new constraints, CHOCOLATE and GUM. The constraint labeled CHOCOLATE forces ICHOCO to equal one when chocolates are manufactured. Similarly, the constraint labeled GUM forces IGUMDR to equal one when gumdrops are manufactured. Also, you should include a constraint labeled ONLY_ONE that requires the sum of ICHOCO and IGUMDR to equal one. (Note that this could be accomplished more simply by including ICHOCO and IGUMDR in a SOSEQ set.) Since ICHOCO and IGUMDR are integer variables, this constraint eliminates the possibility of both products being manufactured. Notice the coefficients -10000, which are used to force ICHOCO or IGUMDR to 1 whenever CHOCO and GUMDR are nonzero. This technique, which is often used in integer programming, can cause severe numerical problems. If this driving coefficient is too large, then arithmetic overflows and underflow may result. If the driving coefficient is too small, then the integer variable may not be driven to one as desired by the modeler.

The objective coefficients of the integer variables ICHOCO and IGUMDR are the negatives of the setup costs for the two products. The following is the data set that describes this problem and the call to PROC LP to solve it:

data;								
input _r	ow_	\$10.	choco	gumdr	ichoco	igumdr	_type_	\$ _rhs_;
dataline	s;							
object		.25	.75	-100	-75	max	•	
cooking		15	40	0	0	le	27000	
color		0	56.25	0	0	le	27000	
package	18.	.75	0	0	0	le	27000	
condiments		12	50	0	0	le	27000	
chocolate		1	0	-10000	0	le	0	
gum		0	1	0	-10000	le	0	
only_one		0	0	1	1	eq	1	
binary		•	•	1	2	binary	•	
;								
proc lp;								
run;								

The solution shows that gumdrops are produced. See Output 3.8.1.

The LP Proce	lure								
Produktion downward									
Problem Summary									
Objective Eurotion	Max object								
Dbg Variable	max object								
Type Variable									
Problem Dengity (%)	25 71								
FIODIEM DENSICY (%)	23./1								
Variables	Number								
Variabies	Mulliper								
Non-negative	2								
Binary	2								
Slack	6								
	-								
Total	10								
Constraints	Number								
LE	6								
EQ	1								
Objective	1								
Total	8								

Output 3.8.1. Summaries and an Integer Programming Iteration Log

The LP Procedure										
			Integer	Iteration	Log					
Iter	Problem	Condition	Objective	Branched	Value	Sinfeas	Active	Proximity		
1	0	ACTIVE	397.5	ichoco	0.1	0.2	2			
2	-1	SUBOPTIMAL	260				1	70		
3	1	SUBOPTIMAL	285	•		•	0	•		

The LP Procedure										
Solution Summary	Solution Summary									
Integer Optimal Solution	Integer Optimal Solution									
Objective Value	285									
Phase 1 Iterations	1									
Phase 2 Iterations	5									
Phase 3 Iterations	5									
Integer Iterations	3									
Integer Solutions	2									
Initial Basic Feasible Variables	8									
Time Used (seconds)	0									
Number of Inversions	5									
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									

The LP Procedure											
Variable Summary											
	Variable Reduced										
	Col	Name	Status	Туре	Price	Activity	Cost				
	1	choco	DEGEN	NON-NEG	0.25	0	0				
	2	gumdr	BASIC	NON-NEG	0.75	480	0				
	3	ichoco		BINARY	-100	0	2475				
	4	igumdr	BASIC	BINARY	-75	1	0				
	5	cooking	BASIC	SLACK	0	7800	0				
	6	color		SLACK	0	0	-0.013333				
	7	package	BASIC	SLACK	0	27000	0				
	8	condiments	BASIC	SLACK	0	3000	0				
	9	chocolate		SLACK	0	0	-0.25				
	10	gum	BASIC	SLACK	0	9520	0				

The LP Procedure										
Constraint Summary										
	Constraint S/S Dual									
Row	Name	Туре	Col	Rhs	Activity	Activity				
1	object	OBJECTVE		0	285	•				
2	cooking	LE	5	27000	19200	0				
3	color	LE	6	27000	27000	0.0133333				
4	package	LE	7	27000	0	0				
5	condiments	LE	8	27000	24000	0				
6	chocolate	LE	9	0	0	0.25				
7	gum	LE	10	0	-9520	0				
8	only_one	EQ	•	1	1	-75				

The branch and bound tree can be reconstructed from the information contained in the integer iteration log. The column labeled lter numbers the integer iterations. The column labeled Problem identifies the lter number of the parent problem from which the current problem is defined. For example, lter=2 has Problem=-1. This means that problem 2 is a direct descendent of problem 1. Furthermore, because problem 1 branched on ICHOCO, you know that problem 2 is identical to problem 1 with an additional constraint on variable ICHOCO. The minus sign in the Problem=-1 in lter=2 tells you that the new constraint on variable ICHOCO is a lower bound. Moreover, because Value=.1 in lter=1, you know that ICHOCO=.1 in lter=1 so that the added constraint in lter=2 is ICHOCO $\leq [.1]$. In this way, the information in the log can be used to reconstruct the branch and bound tree. In fact, when you save an ACTIVEOUT= data set, it contains information in this format that is used to reconstruct the tree when you restart a problem using the ACTIVEIN= data set. See Example 3.10.

Note that if you defined a *SOSEQ* special ordered set containing the variables CHOCO and GUMDR, the integer variable ICHOCO and IGUMDR and the three associated constraints would not have been needed.

Example 3.9. An Infeasible Problem

This is an example of the Infeasible Information Summary that is displayed when an infeasible problem is encountered. Consider the following problem:

```
max x + y + z + w
subject to: x +3y +2z +4w <= 5
3x + y +2z + w <= 4
5x +3y +3z +3w = 9
x, y, z, w >=0
```

Examination of this problem reveals that it is unsolvable. Consequently, PROC LP identifies it as infeasible. The following program attempts to solve it.

data infeas;										
input	_id_	\$6.	x1-x4	_typ	e_\$_	_rhs_;				
datalines;										
profit	1	1	1	1	max	•				
const1	1	3	2	4	le	5				
const2	3	1	2	1	le	4				
const3	5	3	3	3	eq	9				
;										
proc lp;	run;									

The results are shown in Output 3.9.1.

Output 3.9.1. The Solution of an Infeasible Problem

T	ne LP Procedure
I	Problem Summary
Objective Rhs Varial Type Varia Problem De	FunctionMax profitble_rhs_able_type_ensity (%)77.78
Variables	Number
Non-negat: Slack	ive
Total	
Constraint	ts Number
LE EQ Objective	-
Total	

ERROR: Infeasible problem. Note the constraints in the constraint summary that are identified as infeasible. If none of the constraints are flagged then check the implicit bounds on the variables.

The LP Procedure		
Solution Summary		
Infeasible Problem		
Objective Value	2.25	
Phase 1 Iterations	2	
Phase 2 Iterations	0	
Phase 3 Iterations	0	
Integer Iterations	0	
Integer Solutions	0	
Initial Basic Feasible Variables	4	
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	

The LP Procedure											
	Variable Summary										
	Variable					Reduced					
Col	Name	Status	Туре	Price	Activity	Cost					
1	x1	BASIC	NON-NEG	1	0.875	0					
2	x2	BASIC	NON-NEG	1	1.375	0					
3	x 3	ALTER	NON-NEG	1	0	0					
4	x 4		NON-NEG	1	0	-0.25					
5	const1		SLACK	0	0	-0.25					
6	const2		SLACK	0	0	-0.25					

The LP Procedure											
	Constraint Summary										
	Constraint		s/s			Dual					
Row	Name	Туре	Col	Rhs	Activity	Activity					
1	profit	OBJECTVE	•	0	2.25						
2	const1	LE	5	5	5	0.25					
3	const2	LE	6	4	4	0.25					
INF	const3	EQ	•	9	8.5	0					

Infeasible Information Summary Infeasible Row const3 Constraint Activity 8.5 Row Type EQ Rhs Data 9 Variable Coefficient Activity Bound Bound x1 5 0.875 0 INFINITY x2 3 1.375 0 INFINITY x3 3 0 0 INFINITY	The LP Procedure								
Infeasible Row const3 Constraint Activity 8.5 Row Type EQ Rhs Data 9 Variable Coefficient Activity Bound Bound x1 5 0.875 0 INFINITY x2 3 1.375 0 INFINITY x3 3 0 0 INFINITY		Infeasible	Information	Summary					
Lower Upper Variable Coefficient Activity Bound Bound x1 5 0.875 0 INFINITY x2 3 1.375 0 INFINITY x3 3 0 0 INFINITY		Infeasible Constraint Row Type Rhs Data	Row Activity	const3 8.5 EQ 9					
x1 5 0.875 0 INFINITY x2 3 1.375 0 INFINITY x3 3 0 0 INFINITY	Variable	Coefficient	Activity	Lower Bound	Upper Bound				
	x1 x2 x3	5 3 3	0.875 1.375 0	0 0 0	INFINITY INFINITY INFINITY				

Note the information given in the Infeasible Information Summary for the infeasible row CONST3. It shows that the equality constrained row CONST3 with right-hand-side 9 was found to be infeasible with activity 8.5. The summary also shows each variable that has a nonzero coefficient in that row and its activity level at the infeasibility. Examination of these model parameters might give you a clue as to the cause of infeasibility, such as an incorrectly entered coefficient or right-hand-side value.

Example 3.10. Restarting an Integer Program

The following example is attributed to Haldi (Garfinkel and Nemhauser 1972) and is used in the literature as a test problem. Notice that the ACTIVEOUT= and the PRI-MALOUT= options are used when invoking PROC LP. These cause the LP procedure to save the primal solution in the data set named P and the active tree in the data set named A. If the procedure fails to find an optimal integer solution on the initial call, it can be called later using the A and P data sets as starting information.

```
data haldi10;
    input x1-x12 _type_ $ _rhs_;
    datalines;
  0
        0
              0
                    0
                          0
                                0
                                      1
                                           1
                                                 1
                                                       1
                                                             1
                                                                  1
                                                                       MAX
        7
                                5
                                           7
  9
             16
                    8
                        24
                                      3
                                                 8
                                                       4
                                                             6
                                                                   5
                                                                       \mathbf{LE}
                                                                             110
                    2
                        20
                                8
                                      4
                                                             5
                                                                   8
 12
        6
              6
                                           6
                                                 3
                                                       1
                                                                       LE
                                                                               95
 15
        5
             12
                    4
                          4
                                5
                                     5
                                           5
                                                 6
                                                       2
                                                             1
                                                                   5
                                                                       LE
                                                                               80
 18
        4
              4
                   18
                        28
                                1
                                     6
                                           4
                                                 2
                                                       9
                                                             7
                                                                   1
                                                                       \mathbf{LE}
                                                                             100
-12
        0
              0
                    0
                          0
                                0
                                     1
                                           0
                                                 0
                                                       0
                                                             0
                                                                   0
                                                                       \mathbf{LE}
                                                                                0
  0
     -15
              0
                    0
                          0
                                0
                                     0
                                           1
                                                 0
                                                       0
                                                             0
                                                                   0
                                                                       \mathbf{LE}
                                                                                0
  0
        0
           -12
                    0
                          0
                                0
                                      0
                                           0
                                                 1
                                                       0
                                                             0
                                                                   0
                                                                       \mathbf{LE}
                                                                                0
                 -10
                          0
                                0
                                      0
                                           0
                                                 0
                                                       1
  0
        0
              0
                                                             0
                                                                   0
                                                                       LE
                                                                                0
   0
              0
                    0
                      -11
                                0
                                      0
                                           0
                                                 0
                                                       0
                                                             1
                                                                   0
                                                                       \mathbf{LE}
                                                                                0
        0
                    0
                          0
                            -11
                                      0
                                           0
                                                 0
                                                       0
                                                             0
                                                                   1
  0
        0
              0
                                                                       \mathbf{LE}
                                                                                0
  1
        1
              1
                    1
                          1
                                1 1000
                                         1000 1000 1000 1000 1000 UPPERBD .
        2
              3
                    4
                          5
                                6
                                      7
                                           8
                                                     10
                                                                 12
                                                                       INTEGER .
  1
                                                 9
                                                           11
   ;
```

```
proc lp ACTIVEOUT=a PRIMALOUT=p;
run;
```

The ACTIVEOUT= data set contains a representation of the current active problems in the branch and bound tree. The PRIMALOUT= data set contains a representation of the solution to the current problem. These two can be used to restore the procedure to an equivalent state to the one it was in when it stopped.

The results from the call to PROC LP is shown in Output 3.10.1. Notice that the procedure performed 100 iterations and then terminated on maximum integer iterations. This is because, by default, IMAXIT=100. The procedure reports the current best integer solution.

Output 3.10.1. Output from the HALDI10 Problem

```
The LP Procedure
         Problem Summary
                       Max _OBS1_
Objective Function
Rhs Variable
                             _RHS_
Type Variable
                            _TYPE_
Problem Density (%)
                             31.82
Variables
                            Number
Integer
                                  6
Binary
                                  6
Slack
                                 10
                                 22
Total
Constraints
                            Number
\mathbf{LE}
                                 10
Objective
                                 1
Total
                                11
```

The LP Procedure										
			Integer 1	Iteration	Log					
Iter	Problem	Condition	Objective	Branched	Value	Sinfeas	Active	Proximity		
1	0	ACTIVE	18.709524	X9	1.543	1.11905	2	•		
2	-1	ACTIVE	18.676973	X11	0.487	1.96707	3			
3	-2	ACTIVE	18.642641	X9	2.482	1.49839	4	•		
4	-3	ACTIVE	18.402785	X12	7.379	1.62351	5			
5	4	ACTIVE	18.296842	X6	0.636	1.94504	6	•		
6	-5	ACTIVE	18.056766	X4	0.687	0.99534	7	•		
7	6	ACTIVE	15.796152	X12	1.532	2.27511	8	•		
8	7	ACTIVE	15.69596	X1	0.434	2.48667	9			
9	8	ACTIVE	13.547078	X8	8.547	1.22403	10			
10	9	ACTIVE	13.499259	X11	1.499	1.35222	11			
11	-10	ACTIVE	13.451299	X2	0.497	1.37987	12	•		
12	11	ACTIVE	12.8	X3	0.25	0.65	12			
13	-12	ACTIVE	11.454545	X9	8.455	0.63636	13	•		
14	13	ACTIVE	11.444444	X11	2.444	0.66667	14	•		
15	-14	ACTIVE	11.431818	X9	7.432	0.70455	15	•		
16	15	ACTIVE	11.422222	X11	3.422	0.73333	16	•		
17	-16	ACTIVE	11.409091	X9	6.409	0.77273	17	•		
18	17	ACTIVE	11.4	X11	4.4	0.8	18	•		
19	-18	ACTIVE	11.386364	X5	0.455	0.84091	18			
20	-19	SUBOPTIMAL	10				17	8		
21	18	ACTIVE	11	x5	0.5	0.5	18	8		
22	-21	FATHOMED	9.5				17	8		
23	21	FATHOMED	7				16	8		
24	-17	FATHOMED	11.054545				15	8		
25	16	ACTIVE	11	x5	0.417	0.41667	16	8		
26	-25	FATHOMED	9.25				15	8		
27	25	FATHOMED	8				14	8		
28	-15	FATHOMED	11.090909				13	8		
29	14	ACTIVE	11	x5	0.333	0.33333	14	8		
30	-29	FATHOMED	9.8571429				13	8		
31	2.9	FATHOMED	9				12	8		
32	-1.3	FATHOMED	11,127273				11	8		
33	-11	ACTIVE	12,948052	x3	0.25	0.48377	11	8		
34	-33	ACTIVE	11.233766	x8	5.234	0.41558	11	8		
35	34	ACTIVE	11.204545	x9	3.205	0.38636	12	8		
35	35	ACTIVE	11.2	x11	2.2	0.4	13	8		
30	-36	FATHOMED	11.064935		2.2	0.1	12	8		
38	36	ACTIVE		x5	0.25	0.25	13	8		
39	-38	FATHOMED	9,9015152				12	8		
40	38	FATHOMED	9		•	•	11	8		
41	-35	FATHOMED	11,090909	-	•	•	10	8		
42	10	ACTIVE	13,437662	x2	0.533	1.28171	11	8		
43	42	ACTIVE	11.805195	x9	9,805	0.46861	12	8		
44	-43	ACTIVE	11,69697	x12	0.697	0.56061	12	Ř		
45	-44	INFEASTRIE	11,59596		5.557	5.55001	12	8		
46	44	ACTIVE	11.373377	- X9	10.37	0.59984	12	8		
47	-46	INFEASTRIE	10,440404		10.57		12	8		
49	46	ACTIVE	11	• x5	0.236	0.40278	12	8		
40	-49	FATHOMED	9.5714286		5.250	0.102/0	12	s s		
		TNFEASTRIF	10	•	•	•	11	S R		
50	-10	THE BUO TOUR	10	•	•	•		0		

51	43	ACTIVE	11	X3	0.75	0.34091	11	8
52	-51	ACTIVE	11	X5	0.091	0.09091	11	8
53	-52	FATHOMED	9.3181818	•			10	8
54	-42	ACTIVE	13.087662	X3	0.257	0.43588	10	8
55	-54	ACTIVE	11.402597	X8	6.403	0.49351	11	8
56	55	ACTIVE	11.352273	X9	3.352	0.44318	12	8
57	5	ACTIVE	16.338433	X10	3.532	2.59511	13	8
58	57	ACTIVE	16.305108	X11	1.441	2.04051	14	8
59	-58	ACTIVE	16.262829	X2	0.376	2.29921	15	8
60	59	ACTIVE	15.823255	X11	2.504	2.47733	16	8
61	-60	ACTIVE	15.815163	X3	0.522	2.15025	17	8
62	61	INFEASIBLE	13.195402				16	8
63	-61	ACTIVE	14.859188	X1	0.354	1.71629	17	8
64	63	ACTIVE	12.881818	X4	0.3	0.69091	18	8
65	64	ACTIVE	11.681818	X9	8.682	0.59091	19	8
66	65	ACTIVE	11.666667	X5	0.333	0.66667	19	8
67	-66	FATHOMED	10.666667				18	8
68	-65	INFEASIBLE	11.321212				17	8
69	-64	ACTIVE	12.181818	x5	0.273	0.45455	17	8
70	-69	FATHOMED	9.7272727				16	8
70	-63	ACTIVE	13,121891	• x5	0.423	1.24129	16	8
72	71	INFEASTBLE	11.755				15	8
72	60	ACTIVE	15 507892	• x 3	• 0 489	• 1 84958	16	8
75	73	INFEASIBLE	13 444444	10	0.105	1.01950	15	8
75	-73	ACTIVE	14 549752	• x9	• 6 453	• 1 28964	16	8
75	75	ACTIVE	14 531636	x7	3 532	1 24448	17	8
70	-76	ACTIVE	14 512121	x0	5 512	1 30303	18	8
78	-70	ACTIVE	14 401636	x7	4 492	1 34776	19	8
70	-78	ACTIVE	14 354545	x0	4 548	1 52462	20	8
80	70	ACTIVE	14 111111	v1	0 472	1 68182	20	8
81	-80	ACTIVE	12 441212	X0	3 303	0 82788	22	8
01	-00	ACTIVE	12 205675	x7	5 5 9 5	0.02700	22	0
82	- 92	ACIIVE	11 0/607	x f	0 192	0.37921	23	0
0.0	-02	TNEEACTOLE	11 002222	AD .	0.102	0.32955	23	0
04	03	LINFEASIBLE	12 121212	• •	• • •	•	22	0
05	04	ACIIVE		A4	0.212	0.51515	22	0
00	-05	ATHOMED	11 520202	• v10	• • • •		21	0
07	-01	ACTIVE	11 254545	A10 V7	U.55	0.70455	22	0
00	0/	ACTIVE	11.254545	A/	5.255	0.43030	22	0
89	00	ACTIVE	11	X.D	0.182	0.18182	22	8
90	-89	FATHOMED	11 192222	•	•	•	21	8
91	-87	INFEASIBLE	11.1/3333	•	•	•	20	8
92	80	INFEASIBLE	12.380471	•	•	•	19	8
93	-79	FATHOMED	13.0/0136	•		•	18	ŏ
94	78	ACTIVE	14	X1 7710	0.538	0.94364	19	×
95	-94	ACTIVE	12.534545	XT0	2.482	0.96455	20	×
96	95	ACTIVE	12.245455	х9	4.245	0.62727	21	ъ С
97	-96	FATHOMED	11.113636	•	•		20	8
98	96	ACTIVE	12	X4	0.2	0.38182	21	8
99	-98	FATHOMED	T0.989899	•	•	•	20	8
100	98	FATHOMED	10	•	•	•	19	8

WARNING: The maximum number of integer iterations has been exceeded. Increase this limit with the 'IMAXIT=' option on the RESET statement.

The LP Procedure		
Solution Summary		
Terminated on Maximum Integer	Iterations	
Integer Feasible Soluti	on	
Objective Value	10	
Phase 1 Iterations	0	
Phase 2 Iterations	16	
Phase 3 Iterations	240	
Integer Iterations	100	
Integer Solutions	1	
Initial Basic Feasible Variables	12	
Time Used (seconds)	0	
Number of Inversions	37	
The state of the second s	1 - 0	
Epsilon Teficitor	1 707(0)7200	
Infinity Mawimum Dhaga 1 Thomationg	1.00	
Maximum Phase 1 Iterations	100	
Maximum Phase 2 Iterations	001	
Maximum Phase 3 Iterations	100	
Maximum integer iterations	100	
TIME LIMIT (Seconds)	120	

		2	The LP Proce	edure			
		7	/ariable Su	mary			
	Variable					Reduced	
Col	Name	Status	Туре	Price	Activity	Cost	
				-			
1	X1		BINARY	0	0	4.5	
2	X2	DEGEN	BINARY	0	0	0	
3	X3		BINARY	0	1	-2.666667	
4	X4		BINARY	0	0	2	
5	X5		BINARY	0	1	-4	
6	X6		BINARY	0	1	-0.833333	
7	X7	DEGEN	INTEGER	1	0	0	
8	X8		INTEGER	1	0	-0.244444	
9	X9		INTEGER	1	3	-0.333333	
10	X10	DEGEN	INTEGER	1	0	0	
11	X11	BASIC	INTEGER	1	6	0	
12	X12		INTEGER	1	1	0.1666667	
13	_OBS2_		SLACK	0	0	-0.166667	
14	_OBS3_	BASIC	SLACK	0	14	0	
15	_OBS4_	BASIC	SLACK	0	30	0	
16	_OBS5_	BASIC	SLACK	0	18	0	
17	_OBS6_		SLACK	0	0	-0.5	
18	_OBS7_		SLACK	0	0	-0.077778	
19	OBS8	BASIC	SLACK	0	9	0	
20	OBS9		SLACK	0	0	-0.333333	
21	OBS10	BASIC	SLACK	0	5	0	
22	_OBS11_	BASIC	SLACK	0	10	0	

	The	LP Pro	cedure			
	Const	raint	Summary			
Constrai	nt	s/s			Dual	
Row Name	Туре	Col	Rhs	Activity	Activity	
1 _OBS1_	OBJECTVE		0	10	•	
2 _OBS2_	LE	13	110	110	0.1666667	
3 _OBS3_	LE	14	95	81	0	
4 _OBS4_	LE	15	80	50	0	
5 _OBS5_	LE	16	100	82	0	
6 _OBS6_	LE	17	0	0	0.5	
7 _OBS7_	LE	18	0	0	0.0777778	
8 _OBS8_	LE	19	0	-9	0	
9 _OBS9_	LE	20	0	0	0.3333333	
10 _OBS10_	LE	21	0	-5	0	
 11OBS11	LE	22	0	-10	0	

To continue with the solution of this problem, invoke PROC LP with the ACTIVEIN= and PRIMALIN= options and reset the IMAXIT= option. This restores the branch and bound tree and simplifies calculating a basic feasible solution from which to start processing.

proc lp data=haldi10 ACTIVEIN=a PRIMALIN=p IMAXIT=250; run;

The procedure picks up iterating from a equivalent state to where it left off. The problem will still not be solved when IMAXIT=250 occurs.

Example 3.11. Alternative Search of the Branch and Bound Tree

In this example, the HALD110 problem is solved. However, here the default strategy for searching the branch and bound tree is modified. By default, the search strategy has VARSELECT=*FAR*. This means that when searching for an integer variable on which to branch, the procedure uses the one that has a value farthest from an integer value. An alternative strategy has VARSELECT=*PENALTY*. This strategy causes PROC LP to look at the cost, in terms of the objective function, of branching on an integer variable. The procedure looks at PENALTYDEPTH= integer variables before choosing the one with the largest cost. This is a much more expensive strategy (in terms of execution time) than the VARSELECT=*FAR* strategy, but it can be beneficial if fewer integer iterations must be done to find an optimal solution.

proc lp VARSELECT=PENALTY; run;

Compare the number of integer iterations needed to solve the problem using this heuristic with the default strategy used in Example 3.10. In this example, the difference is profound; in general, solution times can vary significantly with the search technique. See Output 3.11.1.

The LP Proced	lure	
Problem Summary		
Objective Function Rhs Variable Type Variable	Max _OBS1_ _RHS_ TYPE	
Problem Density (%)	31.82	
Variables	Number	
Integer Binary	6 6	
Slack	10	
Total	22	
Constraints	Number	
LE	10	
ODJECTIVE	T	
Total	11	

Output 3.11.1. Summaries and an Integer Programming Iteration Log: Using VARSELECT=PENALTY
The LP Procedure								
			Integer 1	Iteration	Log			
Iter	Problem	Condition	Objective	Branched	Value	Sinfeas	Active	Proximity
1	0	ACTIVE	18.709524	X3	0.129	1.11905	2	•
2	-1	ACTIVE	17.063758	X4	0.694	0.89262	3	
3	2	ACTIVE	15.225596	X5	0.459	2.34336	4	
4	3	ACTIVE	12,7375	x2	0.812	0.875	5	
5	-4	ACTIVE	12.55	X6	0.05	0.5	6	
6	5	ACTIVE	12.54	x7	0.24	0.56	7	
7	-6	ACTIVE	12.35	X1	0.083	0.43333	7	
8	-7	ACTIVE	9.6	X8	8.6	0.4	8	
9	8	ACTIVE	9.6	x7	1.6	0.4	9	
10	-9	ACTIVE	9.6	x8	7.6	0.4	10	
11	10	ACTIVE	9.6	X7	2.6	0.4	11	
12	-11	ACTIVE	9.6	x8	6.6	0.4	12	
13	12	ACTIVE	9.6	x7	3.6	0.4	13	
14	-13	ACTIVE	9.6	x8	5.6	0.4	14	
15	14	ACTIVE	9.6	x7	4.6	0.4	15	
16	-15	ACTIVE	9.6	x8	4.6	0.4	16	
17	16	ACTIVE	9.6	x7	5.6	0.4	17	•
18	-17	ACTIVE	9.6	x8	3.6	0.4	18	
19	18	ACTIVE	9.6	x7	6.6	0.4	19	•
20	-19	ACTIVE	9.6	x8	2.6	0.4	20	
21	20	ACTIVE	9.6	x7	7.6	0.4	21	
22	-21	ACTIVE	9.6	x8	1.6	0.4	22	•
23	22	ACTIVE	9.6	x7	8.6	0.4	23	•
24	-23	ACTIVE	9.6	x8	0.6	0.4	24	•
25	24	ACTIVE	9.6	x7	9.6	0.4	25	•
26	-25	INFEASIBLE	8.5333333				24	
27	25	ACTIVE	9.5	x9	0.5	0.5	25	
28	-27	INFEASIBLE	7.8				24	
29	27	SUBOPTIMAL	9				6	9
30	6	ACTIVE	12,428571	x8	12.43	0.42857	7	9
31	30	ACTIVE	12.375	x9	0.375	0.375	8	9
32	31	SUBOPTIMAL	12	•			3	6
33	-3	FATHOMED	13.063239				2	6
34	-2	ACTIVE	16.547134	X6	0.816	0.78733	3	6
35	34	ACTIVE	14.46789	X1	0.128	0.91743	4	6
36	35	ACTIVE	14.28125	X9	7.219	0.28125	5	6
37	36	ACTIVE	14.257979	X8	0.279	0.31915	6	6
38	-37	ACTIVE	14.197917	X2	0.067	0.73958	6	6
39	-38	ACTIVE	13.25	X10	6.5	0.75	7	6
40	39	ACTIVE	13.222222	X9	4.444	0.66667	8	6
41	40	ACTIVE	13.212766	X10	5.83	0.55319	9	6
42	41	ACTIVE	13.166667	X8	6.333	0.5	10	6
43	-42	ACTIVE	13.15625	X10	4.812	0.53125	11	6
44	43	SUBOPTIMAL	13				4	5
45	37	ACTIVE	14.111111	X10	7.111	0.11111	4	5
46	45	ACTIVE	14.104762	X5	0.01	0.11429	5	5
47	-46	FATHOMED	10.230769	•	•		4	5
48	46	SUBOPTIMAL	14	•	•		2	4
49	-34	ACTIVE	16.363499	X12	8.796	0.78733	3	4
50	49	ACTIVE	16.220741	X9	0.782	0.86074	4	4

		,						
!	51 -50	ACTIVE	16.179592	X12	7.776	0.72653	5	4
!	52 51	ACTIVE	15.997778	X9	1.633	1.07778	6	4
!	53 52	ACTIVE	15.927381	X2	0.054	0.77381	7	4
!	54 53	ACTIVE	15.660606	X1	0.053	0.45253	8	4
!	5 -54	FATHOMED	13.386667				7	4
!	6 54	ACTIVE	15.555556	X10	7.556	0.44444	8	4
!	57 56	ACTIVE	15.52381	X5	0.048	0.52381	9	4
!	8 -57	FATHOMED	12.853333	•	•		8	4
!	59 57	SUBOPTIMAL	15	•	•		3	3
	50 50	FATHOMED	16.117029	•	•		2	3
(51 -49	SUBOPTIMAL	16	•	•		1	2
	52 1	ACTIVE	18.0223	X4	0.827	0.81338	2	2
(53 62	FATHOMED	15.927292	•	•	•	1	1
(54 -62	ACTIVE	17.67723	X10	7.93	0.43662	2	1
(5 64	ACTIVE	17.488782	X6	0.737	0.91026	3	1
(6 -65	ACTIVE	17.225962	X8	2.38	0.69231	4	1
(66	ACTIVE	17.221818	X1	0.016	0.37111	5	1
(58 -67	FATHOMED	14.5	•	•	•	4	1
(67 67	FATHOMED	17.172375	•	•	•	3	1
	/0 -66	FATHOMED	17.138028	•	•	•	2	1
	1 65	FATHOMED	15.784375	•	•		1	1
	2 -64	ACTIVE	17.166667	X6	0.833	0.33333	1	1
	/3 -72	SUBOPTIMAL	17	•	•		0	•

The LP Procedure									
Solution Summary									
Integer Optimal Solution									
Objective Value	17								
Phase 1 Iterations	0								
Phase 2 Iterations	16								
Phase 3 Iterations	111								
Integer Iterations	73								
Integer Solutions	7								
Initial Basic Feasible Variables	12								
Time Used (seconds)	0								
Number of Inversions	20								
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								

The LP Procedure										
Variable Summary										
Variab	le				Reduced					
Col Name	Status	Туре	Price	Activity	Cost					
1 X1	DEGEN	BINARY	0	0	0					
2 X2	DEGEN	BINARY	0	0	0					
3 X3	DEGEN	BINARY	0	0	0					
4 X4		BINARY	0	1	-18					
5 X5	DEGEN	BINARY	0	0	0					
6 X6		BINARY	0	1	-1					
7 X7		INTEGER	1	0	-6.5					
8 X8		INTEGER	1	0	-3.266667					
9 X9		INTEGER	1	0	-1.333333					
10 X10		INTEGER	1	8	-8					
11 X11		INTEGER	1	0	-8.545455					
12 X12	BASIC	INTEGER	1	9	0					
13 _OBS2_	BASIC	SLACK	0	20	0					
14 _OBS3_	BASIC	SLACK	0	5	0					
15 _OBS4_	BASIC	SLACK	0	10	0					
16 _OBS5_		SLACK	0	0	-1					
17 _OBS6_		SLACK	0	0	-1.5					
18 _OBS7_		SLACK	0	0	-0.266667					
19 _OBS8_		SLACK	0	0	-0.333333					
20 _OBS9_	BASIC	SLACK	0	2	0					
21 _OBS10	_	SLACK	0	0	-2.545455					
22 _OBS11	BASIC	SLACK	0	2	0					

The LP Procedure											
Constraint Summary											
Co: Row Nat	nstraint me Type	s/s Col	Rhs	Activity	Dual Activity						
1 _0	BS1_ OBJECT	IVE .	0	17							
2 _0	BS2_ LE	13	110	90	0						
3_0	BS3_ LE	14	95	90	0						
4 _0	BS4_ LE	15	80	70	0						
5 _0	BS5_ LE	16	100	100	1						
6_0	BS6_ LE	17	0	0	1.5						
7 _0	BS7_ LE	18	0	0	0.2666667						
8_0	BS8_ LE	19	0	0	0.3333333						
9_0	BS9_ LE	20	0	-2	0						
10 _0	BS10_ LE	21	0	0	2.5454545						
11 _0	BS11_ LE	22	0	-2	0						

Although the VARSELECT=*PENALTY* strategy works well in this example, there is no guarantee that it will work well with your model. Experimentation with various strategies is necessary to find the one that works well with your model and data, particularly if a model is solved repeatedly with few changes to either the structure or the data.

Example 3.12. An Assignment Problem

This example departs somewhat from the emphasis of previous ones. Typically, linear programming models are large, have considerable structure, and are solved with some

164 • Chapter 3. The LP Procedure

regularity. Some form of automatic model building, or matrix generation as it is commonly called, is a useful aid. The sparse input format provides a great deal of flexibility in model specification so that, in many cases, the DATA step can be used to generate the matrix.

The following assignment problem illustrates some techniques in matrix generation. In this example, you have four machines that can produce any of six grades of cloth, and you have five customers that demand various amounts of each grade of cloth. The return from supplying a customer with a demanded grade depends on the machine on which the cloth was made. In addition, the machine capacity depends both upon the specific machine used and the grade of cloth made.

To formulate this problem, let *i* denote customer, *j* denote grade, and *k* denote machine. Then let x_{ijk} denote the amount of cloth grade *j* made on machine *k* for customer *i*; let r_{ijk} denote the return from selling one unit of grade *j* cloth made on machine *k* to customer *i*; let d_{ij} denote the demand for grade *j* cloth by customer *i*; let c_{jk} denote the number of units of machine *k* required to produce one unit of grade *j* cloth; and let a_k denote the number of units of machine *k* available. Then, you get

```
\max \Sigma_{ijk} r_{ijk} x_{ijk}
```

subject to:
$$\begin{split} \Sigma_k x_{ijk} &= d_{ij} & \text{for all } i \text{ and } j \\ \Sigma_{ij} c_{jk} x_{ijk} &\leq a_k & \text{for all } k \\ x_{ijk} &\geq 0 & \text{for all } i, j, \text{ and } k \end{split}$$

The data are saved in three data sets. The OBJECT data set contains the returns for satisfying demand, the DEMAND data set contains the amounts demanded, and the RESOURCE data set contains the conversion factors for each grade and the total amounts of machine resources available.

title 'An Assignment Problem'; data object; input machine customer grade1 grade2 grade3 grade4 grade5 grade6; datalines; 1 1 102 140 105 105 125 148 118 1 2 115 133 118 143 166 1 3 70 108 83 83 88 86 14 79 117 87 87 107 105 1 77 115 90 90 105 148 5 2 1 123 150 125 124 154 2 2 130 157 132 131 166 . 2 3 103 130 115 114 129 • 2 4 101 128 108 107 137 • 2 5 118 145 130 129 154 3 1 83 97 122 147 . . 3 2 119 180 133 163 . 3 3 67 91 101 101 • • 34 85 104 129 • • 129 35 90 • 114 134 179 . 4 1 108 121 79 • 112 132

```
4 2 121 132 92
                        150
                    130
                •
4 3 78 91 59
                    77
                         72
                •
                    109 104
4 4 100 113 76
                •
4 5 96 109 77
                    105 145
;
data demand;
  input customer
        grade1 grade2 grade3 grade4 grade5 grade6;
  datalines;
1 100 100 150 150 175 250
2 300 125 300 275 310
                       325
3 400
       0 400 500 340
                       0
4 250
       0 750 750
                  0
                         0
   0 600 300 0 210 360
5
;
data resource;
  input machine
        grade1 grade2 grade3 grade4 grade5 grade6 avail;
  datalines;
1 .250 .275 .300 .350 .310
                           .295 744
2 .300 .300 .305 .315 .320 .
                                 244
3 .350 . . .320 .315 .300 790
4 .280 .275 .260 .
                      .250 .295 672
;
```

The linear program is built using the DATA step. The model is saved in a SAS data set in the sparse input format for PROC LP. Each section of the following DATA step generates a piece of the linear program. The first section generates the objective function; the next section generates the demand constraints; and the last section generates the machine resource availability constraints.

```
/* build the linear programming model */
data model;
array grade{6} grade1-grade6;
length _type_ $ 8 _row_ $ 8 _col_ $ 8;
keep _type_ _row_ _col_ _coef_;
ncust=5;
nmach=4;
ngrade=6;
/* generate the objective function */
_type_='MAX';
_row_='OBJ';
do k=1 to nmach;
do i=1 to ncust;
_ link readobj; /* read the objective coefficient data */
```

```
do j=1 to ngrade;
            if grade{j}^=. then do;
               _col_='X'||put(i,1.)||put(j,1.)||put(k,1.);
               _coef_=grade{j};
               output;
            end;
         end;
      end;
   end;
   /* generate the demand constraints */
   do i=1 to ncust;
      link readdmd;
                           /* read the demand data */
      do j=1 to ngrade;
         if grade{j}^=. then do;
            _type_='EQ';
            _row_='DEMAND'||put(i,1.)||put(j,1.);
            _col_='_RHS_';
            _coef_=grade{j};
            output;
            _type_=' ';
            do k=1 to nmach;
               _col_='X' | |put(i,1.) | |put(j,1.) | |put(k,1.);
               _coef_=1.0;
               output;
            end;
         end;
      end;
   end;
   /* generate the machine constraints */
   do k=1 to nmach;
                           /* read the machine data */
      link readres;
      _type_='LE';
      _row_='MACHINE'||put(k,1.);
      _col_='_RHS_';
      _coef_=avail;
      output;
      _type_=' ';
      do i=1 to ncust;
         do j=1 to ngrade;
            if grade{j}^=. then do;
               _col_='X'||put(i,1.)||put(j,1.)||put(k,1.);
               _coef_=grade{j};
               output;
               end;
            end;
         end;
      end;
readobj: set object;
return;
```

```
readdmd: set demand;
return;
readres: set resource;
return;
run;
```

With the model built and saved in a data set, it is ready for solution using PROC LP. The following program solves the model and saves the solution in the data set called PRIMAL:

```
/* solve the linear program */
proc lp data=model sparsedata noprint primalout=primal;
run;
```

The following output is produced by PROC LP.

Output 3.12.1. Label ?

```
An Assignment Problem
       The LP Procedure
        Problem Summary
Objective Function
                        Max OBJ
Rhs Variable
                         _RHS_
Type Variable
                         _type_
Problem Density (%)
                          5.31
Variables
                         Number
                            120
Non-negative
Slack
                              4
Total
                            124
Constraints
                         Number
LE
                              4
ΕQ
                             30
Objective
                              1
Total
                             35
```

The LP Procedure									
Solution Summary									
Terminated Successfully									
Objective Value	871426.03763								
Phase 1 Iterations	8								
Phase 2 Iterations	43								
Phase 3 Iterations	0								
Integer Iterations	0								
Integer Solutions	0								
Initial Basic Feasible Variables	26								
Time Used (seconds)	0								
Number of Inversions	3								
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								

The solution is prepared for reporting using the DATA step, and a report is written using PROC TABULATE.

```
/* report the solution */
data solution;
   set primal;
  keep customer grade machine amount;
   if substr(_var_,1,1)='X' then do;
     if _value_^=0 then do;
       customer = substr(_var_,2,1);
       grade = substr(_var_,3,1);
      machine = substr(_var_,4,1);
       amount = _value_;
       output;
     end;
  end;
run;
proc tabulate data=solution;
  class customer grade machine;
   var amount;
   table (machine*customer), (grade*amount);
run;
```

The report shown in Output 3.12.2 gives the assignment of customer, grade of cloth, and machine that maximizes the return and does not violate the machine resource availability.

		 1	2	3	4
		+ amount	amount	amount	amount
		+ Sum	+ Sum	Sum	Sum
machine	customer	++ 			
1	1		100.00	150.00	150.00
	2	•	•	300.00	•
	3	·+	·	256.72	210.31
	4	·	•	750.00	•
	5	·	92.27	•	•
2	3	•	•	143.28	•
	5	·	•	300.00	•
3	2	•	•	•	275.00
	3	·	•	•	289.69
	4	·	•	•	750.00
	5	·	•	•	•
4	1	100.00	•	•	•
	2	300.00	125.00	•	•
	3	400.00	•	•	•
	4	250.00	·	•	•
	5	•+ •	507.73		•

Output 3.12.2. An Assignment Problem



Example 3.13. A Scheduling Problem

Scheduling problems is an application area where techniques in model generation can be valuable. Problems involving scheduling are often solved with integer programming and are similar to assignment problems. In this example, you have eight one-hour time slots in each of five days. You have to assign four people to these time slots so that each slot is covered on every day. You allow the people to specify preference data for each slot on each day. In addition, there are constraints that must be satisfied:

- Each person has some slots for which they are unavailable
- Each person must have either slot 4 or 5 off for lunch
- Each person can work only two time slots in a row
- Each person can work only a specified number of hours in the week.

To formulate this problem, let *i* denote person, *j* denote time slot, and *k* denote day. Then, let $x_{ijk} = 1$ if person *i* is assigned to time slot *j* on day *k*, and 0 otherwise; let p_{ijk} denote the preference of person *i* for slot *j* on day *k*; and let h_i denote the number of hours in a week that person *i* will work. Then, you get max $\Sigma_{ijk} p_{ijk} x_{ijk}$

title 'A Scheduling Problem';

 $\begin{array}{lll} \text{subject to:} & \Sigma_i x_{ijk} = 1 & \text{for all } j \text{ and } k \\ & x_{i4k} + x_{i5k} \leq 1 & \text{for all } i, \text{ and } k \\ & x_{i,\ell,k} + x_{i,\ell+1,k} + x_{i,\ell+2,k} \leq 2 & \text{for all } i, \text{ and } k, \text{ and } \ell = 1, \dots, 6 \\ & \Sigma_{jk} x_{ijk} \leq h_i & \text{for all } i \\ & x_{ijk} = 0 \text{ or } 1 & \text{for all } i, \text{ and } k \text{ such that } p_{ijk} > 0 \\ & \text{otherwise } x_{ijk} = 0 \end{array}$

To solve this problem, create a data set that has the hours and preference data for each individual, time slot, and day. A 10 represents the most desirable time slot, and a 1 represents the least desirable time slot. In addition, a 0 indicates that the time slot is not available.

```
data raw;
   input name $ hour slot mon tue wed thu fri;
   datalines;
      20
           1
                10 10 10 10 10
marc
      20
                 9
                    9
                       9
                          9
                             9
           2
marc
      20
           3
                 8
                    8
                       8
                          8
                             8
marc
                   1
marc
      20
           4
                 1
                      1
                          1
                             1
marc
      20
           5
                 1 1 1
                          1
                             1
      20
           6
                 1 1 1
                          1
                             1
marc
      20
           7
                 1
                   1
                      1
                          1
                             1
marc
                 1 1 1
                          1
                             1
marc
      20
           8
mike
      20
           1
                10
                  98
                          7
                             6
mike
      20
           2
                10
                  98
                          7
                             6
mike
      20
           3
                10
                   9
                      8
                          7
                             6
                10
                   33
mike
      20
                          3
                             3
           4
                 1
                   1 1
                          1
mike
      20
           5
                             1
                   2 3
mike
      20
           6
                 1
                          4
                             5
           7
                 1
                   2
                      3
                          4
                             5
mike
      20
      20
                 1
                   2
                      3
                          4
                             5
mike
           8
bill
      20
           1
                10 10 10 10 10
                    9
                      9
bill
      20
           2
                 9
                          9
                             9
                 8
                   8
                      8
                          8
                             8
bill
      20
           3
                   0
                      0
bill
      20
           4
                 0
                          0
                             0
                 1 1 1
bill
      20
           5
                          1
                             1
bill
      20
           6
                 1
                   1
                      1
                          1
                             1
           7
                 1
                   1 1
                          1
bill
      20
                             1
bill
      20
           8
                 1 1 1
                          1
                             1
                          7
                10
                  98
bob
      20
           1
                             6
bob
      20
           2
                10
                   9
                      8
                          7
                             6
                10 9 8 7
bob
      20
           3
                             6
bob
      20
           4
                10 3 3 3
                             3
                1 1 1 1
           5
                             1
bob
      20
                1 2 3 4
bob
      20
           6
                             5
                 1 2 3 4 5
bob
      20
          7
bob
      20
           8
                 1 2 3 4
                             5
;
```

These data are read by the following DATA step, and an integer program is built to solve the problem. The model is saved in the data set named MODEL. First, the objective function is built using the data saved in the RAW data set. Then, the constraints requiring a person to be working in each time slot are built. Next, the constraints allowing each person time for lunch are added. Then, the constraints restricting people to only two consecutive hours are added. Next, the constraints limiting the time that any one person works in a week are added. Finally, the constraints allowing a person to be assigned only to a time slot for which he is available are added. The code to build each of these constraints follows the formulation closely.

```
data model;
  array week {5} mon tue wed thu fri;
   array hours{4} hours1 hours2 hours3 hours4;
  retain hours1-hours4;
  set raw end=eof;
  length _row_ $ 8 _col_ $ 8 _type_ $ 8;
  keep _type_ _col_ _row_ _coef_;
   if
           name='marc' then i=1;
   else if name='mike' then i=2;
   else if name='bill' then i=3;
  else if name='bob' then i=4;
  hours{i}=hour;
   /* build the objective function */
  do k=1 to 5;
      _col_='x'||put(i,1.)||put(slot,1.)||put(k,1.);
      row ='object';
      _coef_=week{k} * 1000;
     output;
      _row_='upper';
      if week{k}=0 then _coef_=1;
      output;
      _row_='integer';
      _coef_=1;
      output;
   end;
   /* build the rest of the model */
   if eof then do;
      _coef_=.;
     _col_=' ';
      _type_='upper';
      _row_='upper';
      output;
      _type_='max';
```

```
_row_='object';
output;
_type_='int';
_row_='integer';
output;
/* every hour 1 person working */
do j=1 to 8;
   do k=1 to 5;
      _row_='work' | |put(j,1.) | |put(k,1.);
      _type_='eq';
      _col_='_RHS_';
      _coef_=1;
      output;
      _coef_=1;
      _type_=' ';
      do i=1 to 4;
         _col_='x'||put(i,1.)||put(j,1.)||put(k,1.);
         output;
      end;
   end;
end;
/* each person has a lunch */
do i=1 to 4;
   do k=1 to 5;
      _row_='lunch'||put(i,1.)||put(k,1.);
      _type_='le';
      _col_='_RHS_';
      _coef_=1;
      output;
      _coef_=1;
      _type_=' ';
      _col_='x'||put(i,1.)||'4'||put(k,1.);
      output;
      _col_='x'||put(i,1.)||'5'||put(k,1.);
      output;
   end;
end;
/* work at most 2 slots in a row */
do i=1 to 4;
   do k=1 to 5;
      do l=1 to 6;
      _row_='seq'||put(i,1.)||put(k,1.)||put(1,1.);
      _type_='le';
      _col_='_RHS_';
      _coef_=2;
       output;
      _coef_=1;
```

```
_type_=' ';
                do j=0 to 2;
                  _col_='x'||put(i,1.)||put(l+j,1.)||put(k,1.);
                  output;
                end;
            end;
         end;
      end;
      /* work at most n hours in a week */
      do i=1 to 4;
         _row_='capacit'||put(i,1.);
         _type_='le';
         _col_='_RHS_';
         _coef_=hours{i};
         output;
         _coef_=1;
         _type_=' ';
         do j=1 to 8;
            do k=1 to 5;
                _col_='x'||put(i,1.)||put(j,1.)||put(k,1.);
               output;
            end;
         end;
      end;
   end;
run;
```

The model saved in the data set named MODEL is in the sparse format. The constraint that requires one person to work in time slot 1 on day 2 is named WORK12; it is $\Sigma_i x_{i12} = 1$.

The following model is saved

TYPE	_COL_	_ROW_	_COEF_
eq	_RHS_	work12	1
	x112	work12	1
	x212	work12	1
	x312	work12	1
	x412	work12	1

in the MODEL data set (which has 1387 observations).

The model is solved using the LP procedure. The option PRIMALOUT=SOLUTION causes PROC LP to save the primal solution in the data set named SOLUTION.

```
/* solve the linear program */
proc lp sparsedata noprint PRIMALOUT=solution
   time=1000 maxit1=1000 maxit2=1000;
run;
```

The following DATA step below takes the solution data set SOLUTION and generates a report data set named REPORT. It translates the variable names x_{ijk} so that a more meaningful report can be written. Then, the PROC TABULATE procedure is used to display a schedule showing how the eight time slots are covered for the week.

```
/* report the solution */
title 'Reported Solution';
data report;
   set solution;
   keep name slot mon tue wed thu fri;
   if substr(_var_,1,1)='x' then do;
     if _value_>0 then do;
        n=substr(_var_,2,1);
        slot=substr(_var_,3,1);
        d=substr(_var_,4,1);
        if
                n='1' then name='marc';
        else if n='2' then name='mike';
        else if n='3' then name='bill';
        else
                           name='bob';
        if
                d='1' then mon=1;
        else if d='2' then tue=1;
        else if d='3' then wed=1;
        else if d='4' then thu=1;
        else
                           fri=1;
        output;
     end;
  end;
run;
proc format;
   value xfmt 1=' xxx
                         ';
run;
proc tabulate data=report;
  class name slot;
   var mon--fri;
   table (slot * name), (mon tue wed thu fri)*sum=' '*f=xfmt.
          /misstext=' ';
run;
```

Output 3.13.1 from PROC TABULATE summarizes the schedule. Notice that the constraint requiring that a person be assigned to each possible time slot on each day is satisfied.

		mon	tue	wed	thu	fri				
slot	name									
1	bill				xxx	xxx				
	marc	xxx	xxx	xxx	ļ	ļ				
2	bob	xxx								
	marc	ļ	xxx	xxx	xxx	xxx				
3	marc	ļ			xxx	xxx				
	mike	xxx	xxx	xxx	ļ					
4	mike	xxx	xxx	xxx	xxx	xxx				
5	marc	xxx	xxx	xxx	xxx	xxx				
6	bob	İ	ļ	xxx	xxx	ĺ				
	mike	xxx	xxx	ļ	ļ	xxx				
7	bob	İ	xxx	İ	İ	ļ				
	marc	xxx			ļ	ļ				
	mike	İ	ļ	xxx	xxx	xxx				
8	bob	ļ				xxx				
	marc	xxx								
	mike		xxx	xxx	xxx					

Output 3.13.1. A Scheduling Problem

Recall that PROC LP puts a character string in the macro variable _ORLP_ that describes the characteristics of the solution on termination. This string can be parsed using macro functions and the information obtained can be used in report writing. The variable can be displayed on the log with

%put &_orlp_;

That statement produces Output 3.13.2.

Output 3.13.2. _ORLP_

```
STATUS=SUCCESSFUL PHASE=3 OBJECTIVE=211000 P_FEAS=YES D_FEAS=YES
INT_ITER=0 INT_FEAS=1 ACTIVE=0 INT_BEST=211000 PHASE1_ITER=49
PHASE2_ITER=55 PHASE3_ITER=0
```

From this you learn, for example, that at termination the solution is integer optimal and has an objective value of 211000.

Example 3.14. A Multicommodity Transshipment Problem with Fixed Charges

The following example illustrates a DATA step program for generating a linear program to solve a multicommodity network flow model that has fixed charges. Consider a network consisting of the following nodes: farm-a, farm-b, farm-c, Chicago, St. Louis, and New York. You can ship three commodities from each farm to Chicago or St. Louis and from Chicago or St. Louis to New York. The following table shows the unit shipping cost for each of the four commodities across each of the arcs. The table also shows the supply (positive numbers) at each of the from nodes and the demand (negative numbers) at each of the to nodes. The fixed charge is a fixed cost for shipping any nonzero amount across an arc. For example, if any amount of any of the four commodities is sent from farm-c to St. Louis, then a fixed charge of 75 units is added to the shipping cost.

 Table 3.8.
 Farms to cities network problem

		Unit	: Sł	nip	ping	Supp	ply an	nd D	emand	Fixed
From	То		Cos	st						Charge
Node	Node	1	2	3	4	1	2	3	4	
farm-a	Chicago	20	15	17	22	100	100	40	•	100
farm-b	Chicago	15	15	15	30	100	200	50	50	75
farm-c	Chicago	30	30	10	10	40	100	75	100	100
farm-a	StLouis	30	25	27	22	•	•	•	•	150
farm-c	StLouis	10	9	11	10	•	•	•	•	75
Chicago	NY	75	75	75	75	-150	-200	-50	-75	200
StLouis	NY	80	80	80	80	•	•	•	•	200

The following program is designed to take the data in the form given in the preceding table. It builds the node arc incidence matrix for a network given in this form and adds integer variables to capture the fixed charge using the type of constraints discussed in Example 3.8. The program solves the model using PROC LP, saves the solution in the PRIMALOUT= data set named SOLUTION, and displays the solution. The DATA step can be easily modified to handle larger problems with similar structure.

title 'Multi-commodity Transshipment Problem with Fixed Charges';

```
data network;
  retain M 1.0e6;
  length _col_ $ 22 _row_ $ 22;
  keep _type_ _col_ _row_ _coef_;
  array sd sd1-sd4;
  array c c1-c4;
  input arc $10. from $ to $ c1 c2 c3 c4 sd1 sd2 sd3 sd4 fx;
  /* for the first observation define some of the rows */
  if _n_=1 then do;
```

```
_type_='upperbd';
   _row_='upper';
   output;
   _type_='lowerbd';
   _row_='lower';
   output;
   _type_='min';
   _row_='obj';
   output;
   _type_='integer';
   _row_='int';
   output;
   end;
_col_='_rhs_';
_type_='le';
                                                                    */
do over sd;
                                    /* loop for each commodity
   coef =sd;
                                    /* the node is a supply node */
   if sd>0 then do;
      _row_=from||' commodity'||put(_i_,2.);
      if from^=' ' then output;
   end:
   else if sd<0 then do;
                                    /* the node is a demand node */
      _row_=to||' commodity'||put(_i_,2.);
      if to<sup>^</sup>=' ' then output;
   end;
   else if from^=' ' & to^=' ' then do; /* a transshipment node */
      _coef_=0;
      _row_=from||' commodity'||put(_i_,2.);
      output;
      _row_=to ||' commodity'||put(_i_,2.);
      output;
   end;
end;
                                    /* loop for each commodity */
do over c;
_col_=arc||' commodity'||put(_i_,2.);
if from<sup>^</sup>=' ' & to<sup>^</sup>=' ' then do;
                                /* add node arc incidence matrix*/
  _type_='le';
  _row_=from || ' commodity' || put(_i_,2.);
  _coef_=1;
 output;
            ||' commodity'||put(_i_,2.);
  _row_=to
  _coef_=-1;
 output;
  _type_=' ';
  _row_='obj';
  _coef_=c;
  output;
                                /* add fixed charge variables
                                                                    */
  _type_='le';
  _row_=arc;
```

```
_coef_=1;output;
     _col_='_rhs_';
    _type_=' ';
     _coef_=0;
    output;
    _col_=arc||'fx';
    _coef_=-M;
    output;
     _row_='int';
    _coef_=1;
    output;
    _row_='obj';
     coef =fx;
    output;
    row ='upper';
     _coef_=1;
    output;
  end;
   end;
   datalines;
a-Chicago farm-a Chicago 20 15 17 22 100 100 40
                                                     . 100
b-Chicago farm-b Chicago 15 15 15 30 100 200 50 50
                                                        75
c-Chicago farm-c Chicago 30 30 10 10
                                      40 100 75 100 100
a-StLouis farm-a StLouis 30 25 27 22
                                                     . 150
                                             •
                                                 •
                                        •
c-StLouis farm-c StLouis 10 9 11 10
                                                        75
                                        •
                                             •
                                                •
                                                    .
Chicago-NY Chicago NY 75 75 75 75 -150 -200 -50 -75 200
StLous-NY StLouis NY
                        80 80 80 80
                                      . . . . . 200
;
/* solve the model */
proc lp sparsedata pout=solution noprint;
run;
/* print the solution */
data;
   set solution;
  rename _var_=arc _value_=amount;
  if _value_^=0 & _type_='NON-NEG';
run;
proc print;
   id arc;
   var amount;
run;
```

The results from this example are shown in Output 3.14.1. The NOPRINT option in the PROC LP statement suppresses the Variable and Constraint Summary sections. This is useful when solving large models for which a report program is available.

179

Here, the solution is saved in data set SOLUTION and reported using PROC PRINT. The solution shows the amount that is shipped over each arc.

Output 3.14.1. Multicommodity Transshipment Problem with Fixed Charges

Multi-commodity Transl	hipment Pro	blem	with F	ixed-Charges
a	rc		amount	
a-Chicago	commodity	1	10	
a-Chicago	commodity	2	100	
b-Chicago	commodity	1	100	
c-Chicago	commodity	3	50	
c-Chicago	commodity	4	75	
c-StLouis	commodity	1	40	
c-StLouis	commodity	2	100	
Chicago-NY	commodity	1	110	
Chicago-NY	commodity	2	100	
Chicago-NY	commodity	3	50	
Chicago-NY	commodity	4	75	
StLous-NY	commodity	1	40	
StLous-NY	commodity	2	100	

References

- Bartels, R. (1971), "A Stabilization of the Simplex Method," *Numerical Mathematics*, 16, 414–434.
- Bland, R.G. (1977), "New Finite Pivoting Rules for the Simplex Method," *Mathematics of Operations Research*, 2, 103–107.
- Breau, R. and Burdet, C.A. (1974), "Branch and Bound Experiments in Zero-one Programming," *Mathematical Programming Study*, ed. M.L. Balinski, 2, 1–50.
- Crowder, H., Johnson, E.L., and Padberg, M.W. (1983), "Solving Large-Scale Zero-One Linear Programming Problems," *Operations Research*, 31, 803–834.
- Dantzig, G.B. (1963), *Linear Programming and Extensions*, Princeton: Princeton University Press.
- Garfinkel, R.S. and Nemhauser, G.L. (1972), *Integer Programming*, New York: John Wiley & Sons.
- Greenberg, H.J. ed. (1978), "Pivot Selection Tactics," in *Design and Implementation* of Optimization Software, Netherlands: Sijthoff & Noordhoff.
- Hadley, G. (1962), Linear Programming, Reading, Massachusetts: Addison-Wesley.
- Harris, P. (1975), "Pivot Selection Methods of the Devex LP Code," in *Mathematical Programming Study 4*, Amsterdam: North-Holland Publishing Co.
- Ignizio, J.P. (1976), *Goal Programming and Extensions*, Lexington, MA: D.C. Heath and Company.
- Murtagh, B.A.(1981) Advanced Linear Programming, Computation and Practice, McGraw-Hill Inc.

Nelson, Mark. (1992), The Data Compression Book, M&T Books.

- Reid, J.K. (1976), "A Sparsity-Exploiting Variant of the Bartels-Golub Decomposition for Linear Programming Bases," *Harwell Report CSS 20*, A.E.R.E., Didcot, Oxfordshire.
- Reid, J.K. (1976), "Fortran Subroutines for Handling Sparse Linear Programming Bases," *Harwell Report AERE-R* 8269, A.E.R.E., Didcot, Oxfordshire.
- Savelsbergh, M.W.P. (1996) "Preprocessing and Probing Techniques for Mixed Integer Programming Problems," ORSA J. on Computing 6, 445–454.

Taha, H.A. (1975), Integer Programming, New York: Academic Press.

The correct bibliographic citation for this manual is as follows: SAS Institute Inc., $SAS/OR^{@}$ User's Guide: Mathematical Programming, Version 8, Cary, NC: SAS Institute Inc., 1999. 566 pp.

SAS/OR[®] 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^{\circledast} 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. $^{\circledast}$ indicates USA registration.

 $IBM^{\circledast}, ACF/VTAM^{\circledast}, AIX^{\circledast}, APPN^{\circledast}, MVS/ESA^{\circledast}, OS/2^{\circledast}, OS/390^{\circledast}, VM/ESA^{\circledast}, and VTAM^{\circledast} are registered trademarks or trademarks of International Business Machines Corporation.$ $<math display="inline">^{\circledast}$ 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.