

Getting Started with Coopr 3.2

COLLABORATORS

	<i>TITLE :</i> Getting Started with Coopr 3.2		
<i>ACTION</i>	<i>NAME</i>	<i>DATE</i>	<i>SIGNATURE</i>
WRITTEN BY	William E. Hart and David L. Woodruff	October 18, 2012	

Contents

1	Release Notes	1
2	Pyomo Overview	2
2.1	Mathematical Modeling	2
2.2	Overview of Modeling Components and Processes	3
2.3	Abstract Versus Concrete Models	4
2.4	A Simple Abstract Pyomo Model	4
2.5	A Simple Concrete Pyomo Model	9
2.6	Solving the Simple Examples	9
3	Sets	11
3.1	Declaration	11
3.2	Operations	12
3.3	Predefined Virtual Sets	13
3.4	Sparse Index Sets	13
3.4.1	Sparse Index Sets Example	16
4	Parameters	18
5	Variables	20
6	Objectives	21
7	Constraints	22
8	Expressions	23
8.1	Rules to Generate Expressions	23
8.2	Piecewise Linear Expressions	24
9	Data Input	27
9.1	Data Command Files	27
9.2	DataPortal Objects	28
9.2.1	Loading Data	28

10 The pyomo Command	31
10.1 Passing Options to a Solver	31
10.2 Troubleshooting	31
10.3 Direct Interfaces to Solvers	32
11 PySP Overview	33
11.1 Overview of Modeling Components and Processes	33
11.2 Birge and Louveaux’s Farmer Problem	33
11.2.1 ReferenceModel.py	34
11.2.2 ReferenceModel.dat	35
11.2.3 ScenarioStructure.dat	36
11.2.4 Scenario data specification	38
11.3 Finding Solutions for Stochastic Models	39
11.3.1 runef	39
11.3.2 runph	39
11.3.3 Solution Output Control	40
11.4 Summary of PySP File Names	40
11.5 Solving Sub-problems in Parallel and/or Remotely	41
12 Scripts	42
12.1 Python Scripts	42
12.1.1 Iterative Example	42
12.2 Changing the Model or Data and Re-solving	45
12.3 Pyomo Callbacks	45
12.3.1 pyomo_preprocess	46
12.3.2 pyomo_create_model	46
12.3.3 pyomo_create_modeldata	46
12.3.4 pyomo_print_model	46
12.3.5 pyomo_modify_instance	47
12.3.6 pyomo_print_instance	47
12.3.7 pyomo_save_instance	47
12.3.8 pyomo_print_results	47
12.3.9 pyomo_save_results	47
12.3.10 pyomo_postprocess	47
12.4 Accessing Variable Values	47
12.4.1 Primal Variable Values	47
12.4.2 One Variable from a Python Script	47
12.4.3 All Variables from a Python Script	48
12.4.4 All Variables from Workflow Callbacks	48

12.5	Accessing Parameter Values	49
12.6	Accessing Duals	49
12.6.1	Access Duals in a Python Script	49
12.6.2	All Duals from Workflow Callbacks	50
12.7	Accessing Solver Status	50
12.8	Display of Solver Output	51
13	Coopr Solver Interfaces	52
14	Using Black-Box Optimizers with Coopr.Opt	53
14.1	Defining and Optimizing Simple Black-Box Applications	53
14.1.1	Defining an Optimization Problem	53
14.1.2	Optimizing with Coliny Solvers	54
14.2	Diving Deeper	56
15	Distributed Optimization with Pyro	59
15.1	Step 1: Starting a Name Server	59
15.2	Step 2: Starting a Dispatch Server	59
15.3	Step 3: Starting a MIP server	59
15.4	Step 4: Running a Client	60
15.5	Moving from Multi-Core to Distributed Computation	60
15.6	Cleaning Up After Yourself	60

Preface

This book provides a quick introduction to Coopr. Coopr is a collection of Python software packages that supports a diverse set of optimization capabilities for formulating and analyzing optimization models. A central component of Coopr is Pyomo, which supports the formulation and analysis of mathematical models for complex optimization applications. This capability is commonly associated with algebraic modeling languages (AMLs), which support the description and analysis of mathematical models with a high-level language. Although most AMLs are implemented in custom modeling languages, Pyomo's modeling objects are embedded within Python, a full-featured high-level programming language that contains a rich set of supporting libraries.

Coopr has also proven an effective framework for developing high-level optimization and analysis tools. For example, the PySP package provides generic solvers for stochastic programming. PySP leverages the fact that Pyomo's modeling objects are embedded within a full-featured high-level programming language, which allows for transparent parallelization of subproblems using Python parallel communication libraries.

Goals of the Book

In this book, we provide a broad overview of different components of the Coopr software. There are roughly two main goals for this book:

1. Help users get started with different Coopr capabilities: Our goal is not to provide a comprehensive reference, but rather to provide a tutorial with simple and illustrative examples. Also, we aim to provide explanations behind the design and philosophy of Coopr.
2. Provide preliminary documentation of new features and capabilities in Coopr. We know that a new feature or capability probably will not be used unless it is documented. As Coopr evolves, we plan to use this book to document these features. This provides users some context concerning the focus of Coopr development, and it also provides an opportunity to get early feedback on new features before they are documented in other contexts.

Who Should Read This Book

This book is intended to be a reference for students, academic researchers and practitioners. Coopr has been effectively used in the classroom with undergraduate and graduate students. However, we assume that the reader is generally familiar with optimization and mathematical modeling. Although this book does not contain a glossary, we recommend the Mathematical Programming Glossary [MPG] as a reference for the reader. We also assume that the reader is generally familiar with the Python programming language. There are a variety of books describing Python, as well as excellent documentation of the Python language and the software packages bundled with Python distributions.

Comments and Questions

Further information about Pyomo and Coopr is available on the Coopr wiki:

<https://software.sandia.gov/trac/coopr>

Coopr is also hosted at COIN-OR:

<https://projects.coin-or.org/Coopr>

We strongly encourage feedback from readers about the software on the Coopr Forum:

coopr-forum@googlegroups.com

If you have comments about this pre-alpha rough draft of the book, please send them to DLWoodruff@UCDavis.edu. We hope this will include feedback on typos and errors in our examples. Additionally, we welcome comments on the presentation of this material, and suggestions for material that we should develop in the other book chapters.

Good Luck!

Chapter 1

Release Notes

The following are highlights of Coopr 3.2:

- Solvers
 - Updates for CBC, Cplex and Gurobi solvers
- Modeling
 - Immutable parameters are now the default
- Other
 - MS Windows installers fixed to work on both Python 2.6 and 2.7
 - MS Windows installers no longer modify the PATH environment
 - Efficiency improvements in model generation, memory, runtime, etc.
 - Restructuring of Pyomo core codes
 - Many bug fixes

The use of immutable parameters may lead to fundamental changes in model interpretation. We now recommend using the following import when defining Coopr models:

```
from __future__ import division
```

This import ensures that integer (or long) parameter values can be used in division in a normal manner. In Python versions before 3.0, the division operator returns the floor of the mathematical result of division if the arguments are integers or longs.

Note

This import line must be at the top of the Pyomo model, before other import lines.



Warning

due to this change, models developed for earlier versions of Coopr may fail to generate correct models with Coopr 3.2. In most cases, including this import statement will resolve this issue.

Chapter 2

Pyomo Overview

2.1 Mathematical Modeling

This chapter provides an introduction to Pyomo: Python Optimization Modeling Objects. A more complete description is contained in [\[PyomoBook\]](#). Pyomo supports the formulation and analysis of mathematical models for complex optimization applications. This capability is commonly associated with algebraic modeling languages (AMLs) such as AMPL [\[AMPL\]](#) AIMMS [\[AIMMS\]](#) and GAMS [\[GAMS\]](#). Pyomo's modeling objects are embedded within Python, a full-featured high-level programming language that contains a rich set of supporting libraries.

Modeling is a fundamental process in many aspects of scientific research, engineering and business. Modeling involves the formulation of a simplified representation of a system or real-world object. Thus, modeling tools like Pyomo can be used in a variety of ways:

- **Explain phenomena** that arise in a system,
- **Make predictions** about future states of a system,
- **Assess key factors** that influence phenomena in a system,
- **Identify extreme states** in a system, that might represent worst-case scenarios or minimal cost plans, and
- **Analyze trade-offs** to support human decision makers.

Mathematical models represent system knowledge with a formalized mathematical language. The following mathematical concepts are central to modern modeling activities:

variables

Variables represent unknown or changing parts of a model (e.g. whether or not to make a decision, or the characteristic of a system outcome). The values taken by the variables are often referred to as a *solution* and are usually an output of the optimization process.

parameters

Parameters represents the data that must be supplied to perform the optimization. In fact, in some settings the word *data* is used in place of the word *parameters*.

relations

These are equations, inequalities or other mathematical relationships that define how different parts of a model are connected to each other.

goals

These are functions that reflect goals and objectives for the system being modeled.

The widespread availability of computing resources has made the numerical analysis of mathematical models a commonplace activity. Without a modeling language, the process of setting up input files, executing a solver and extracting the final results from the solver output is tedious and error prone. This difficulty is compounded in complex, large-scale real-world applications which are difficult to debug when errors occur. Additionally, there are many different formats used by optimization software packages, and few formats are recognized by many optimizers. Thus the application of multiple optimization solvers to analyze a model introduces additional complexities.

Pyomo is an AML that extends Python to include objects for mathematical modeling. Hart et al. [[PyomoBook](#)], [[PyomoJournal](#)] compare Pyomo with other AMLs. Although many good AMLs have been developed for optimization models, the following are motivating factors for the development of Pyomo:

Open Source

Pyomo is developed within Coopr's open source project to promote transparency of the modeling framework and encourage community development of Pyomo capabilities.

Customizable Capability

Pyomo supports a customizable capability through the extensive use of plug-ins to modularize software components.

Solver Integration

Pyomo models can be optimized with solvers that are written either in Python or in compiled, low-level languages.

Programming Language

Pyomo leverages a high-level programming language, which has several advantages over custom AMLs: a very robust language, extensive documentation, a rich set of standard libraries, support for modern programming features like classes and functions, and portability to many platforms.

2.2 Overview of Modeling Components and Processes

Pyomo supports an object-oriented design for the definition of optimization models. The basic steps of a simple modeling process are:

- Create model and declare components
- Instantiate the model
- Apply solver
- Interrogate solver results

In practice, these steps may be applied repeatedly with different data or with different constraints applied to the model. However, we focus on this simple modeling process to illustrate different strategies for modeling with Pyomo.

A Pyomo *model* consists of a collection of modeling *components* that define different aspects of the model. Pyomo includes the modeling components that are commonly supported by modern AMLs: index sets, symbolic parameters, decision variables, objectives, and constraints. These modeling components are defined in Pyomo through the following Python classes:

Set

set data that is used to define a model instance

Param

parameter data that is used to define a model instance

Var

decision variables in a model

Objective

expressions that are minimized or maximized in a model

Constraint

constraint expressions that impose restrictions on variable values in a model

2.3 Abstract Versus Concrete Models

A mathematical model can be defined using symbols that represent data values. For example, the following equations represent a linear program (LP) to find optimal values for the vector x with parameters n and b , and parameter vectors a and c :

$$\begin{aligned} \min \quad & \sum_{j=1}^n c_j x_j \\ \text{s.t.} \quad & \sum_{j=1}^n a_{ij} x_j \geq b_i \quad \forall i = 1 \dots m \\ & x_j \geq 0 \quad \forall j = 1 \dots n \end{aligned}$$

Note

As a convenience, we use the symbol \forall to mean “for all” or “for each.”

We call this an *abstract* or *symbolic* mathematical model since it relies on unspecified parameter values. Data values can be used to specify a *model instance*. The `AbstractModel` class provides a context for defining and initializing abstract optimization models in Pyomo when the data values will be supplied at the time a solution is to be obtained.

In some contexts a mathematical model can be directly defined with the data values supplied at the time of the model definition and built into the model. We call these *concrete* mathematical models. For example, the following LP model is a concrete instance of the previous abstract model:

$$\begin{aligned} \min \quad & 2x_1 + 3x_2 \\ \text{s.t.} \quad & 3x_1 + 4x_2 \geq 1 \\ & x_1, x_2 \geq 0 \end{aligned}$$

The `ConcreteModel` class is used to define concrete optimization models in Pyomo.

2.4 A Simple Abstract Pyomo Model

We repeat the abstract model already given:

$$\begin{aligned} \min \quad & \sum_{j=1}^n c_j x_j \\ \text{s.t.} \quad & \sum_{j=1}^n a_{ij} x_j \geq b_i \quad \forall i = 1 \dots m \\ & x_j \geq 0 \quad \forall j = 1 \dots n \end{aligned}$$

One way to implement this in Pyomo is as follows:

```
from __future__ import division
from coopr.pyomo import *

model = AbstractModel()

model.m = Param(within=NonNegativeIntegers)
model.n = Param(within=NonNegativeIntegers)

model.I = RangeSet(1, model.m)
model.J = RangeSet(1, model.n)

model.a = Param(model.I, model.J)
model.b = Param(model.I)
model.c = Param(model.J)

# the next line declares a variable indexed by the set J
model.x = Var(model.J, domain=NonNegativeReals)

def obj_expression(model):
    return summation(model.c, model.x)

model.OBJ = Objective(rule=obj_expression)
```

```
def ax_constraint_rule(model, i):
    # return the expression for the constraint for i
    return sum(model.a[i,j] * model.x[j] for j in model.J) >= model.b[i]

# the next line creates one constraint for each member of the set model.I
model.AxbConstraint = Constraint(model.I, rule=ax_constraint_rule)
```

Note

Python is interpreted one line at a time. A line continuation character, backslash, is used for Python statements that need to span multiple lines. In Python, indentation has meaning and must be consistent. For example, lines inside a function definition must be indented and the end of the indentation is used by Python to signal the end of the definition.

We will now examine the lines in this example. The first import line is used to ensure that `int` or `long` division arguments are converted to floating point values before division is performed.

```
from __future__ import division
```

In Python versions before 3.0, division returns the floor of the mathematical result of division if arguments are `int` or `long`. This import line avoids unexpected behavior when developing mathematical models with integer values.

The next import line that is required in every Pyomo model. Its purpose is to make the symbols used by Pyomo known to Python.

```
from coopr.pyomo import *
```

The declaration of a model is also required. The use of the name `model` is not required. Almost any name could be used, but we will use the name `model` most of the time in this book. In this example, we are declaring that it will be an abstract model.

```
model = AbstractModel()
```

We declare the parameters `m` and `n` using the Pyomo `Param` function. This function can take a variety of arguments; this example illustrates use of the `within` option that is used by Pyomo to validate the data value that is assigned to the parameter. If this option were not given, then Pyomo would not object to any type of data being assigned to these parameters. As it is, assignment of a value that is not a non-negative integer will result in an error.

```
model.m = Param(within=NonNegativeIntegers)
model.n = Param(within=NonNegativeIntegers)
```

Although not required, it is convenient to define index sets. In this example we use the `RangeSet` function to declare that the sets will be a sequence of integers starting at 1 and ending at a value specified by the parameters `model.m` and `model.n`.

```
model.I = RangeSet(1, model.m)
model.J = RangeSet(1, model.n)
```

The coefficient and right-hand-side data are defined as indexed parameters. When sets are given as arguments to the `Param` function, they indicate that the set will index the parameter.

```
model.a = Param(model.I, model.J)
model.b = Param(model.I)
model.c = Param(model.J)
```

Note

In Python, and therefore in Pyomo, any text after pound sign is considered to be a comment.

The next line interpreted by Python as part of the model declares the variable x . The first argument to the `Var` function is a set, so it is defined as an index set for the variable. In this case the variable has only one index set, but multiple sets could be used as was the case for the declaration of the parameter `model.a`. The second argument specifies a domain for the variable. This information is part of the model and will be passed to the solver when data is provided and the model is solved. Specification of the `NonNegativeReals` domain implements the requirement that the variables be greater than or equal to zero.

```
# the next line declares a variable indexed by the set J
model.x = Var(model.J, domain=NonNegativeReals)
```

In abstract models, Pyomo expressions are usually provided to objective function and constraint declarations via a function defined with a Python `def` statement. The `def` statement establishes a name for a function along with its arguments. When Pyomo uses a function to get objective function or constraint expressions, it always passes in the model (i.e., itself) as the first argument so the model is always the first formal argument when declaring such functions in Pyomo. Additional arguments, if needed, follow. Since summation is an extremely common part of optimization models, Pyomo provides a flexible function to accommodate it. When given two arguments, the `summation` function returns an expression for the sum of the product of the two arguments over their indexes. This only works, of course, if the two arguments have the same indexes. If it is given only one argument it returns an expression for the sum over all indexes of that argument. So in this example, when `summation` is passed the arguments `model.c`, `model.x` it returns an internal representation of the expression $\sum_{j=1}^n c_j x_j$.

```
def obj_expression(model):
    return summation(model.c, model.x)
```

To declare an objective function, the Pyomo function called `Objective` is used. The `rule` argument gives the name of a function that returns the expression to be used. The default `sense` is minimization. For maximization, the `sense=maximize` argument must be used. The name that is declared, which is `OBJ` in this case, appears in some reports and can be almost any name.

```
model.OBJ = Objective(rule=obj_expression)
```

Declaration of constraints is similar. A function is declared to deliver the constraint expression. In this case, there can be multiple constraints of the same form because we index the constraints by i in the expression $\sum_{j=1}^n a_{ij} x_j \geq b_i \quad \forall i = 1 \dots m$, which states that we need a constraint for each value of i from one to m . In order to parametrize the expression by i we include it as a formal parameter to the function that declares the constraint expression. Technically, we could have used anything for this argument, but that might be confusing. Using an `i` for an i seems sensible in this situation.

```
def ax_constraint_rule(model, i):
    # return the expression for the constraint for i
    return sum(model.a[i, j] * model.x[j] for j in model.J) >= model.b[i]
```

Note

In Python, indexes are in square brackets and function arguments are in parentheses.

In order to declare constraints that use this expression, we use the Pyomo `Constraint` function that takes a variety of arguments. In this case, our model specifies that we can have more than one constraint of the same form and we have created a set, `model.I`, over which these constraints can be indexed so that is the first argument to the constraint declaration function. The next argument gives the rule that will be used to generate expressions for the constraints. Taken as a whole, this constraint declaration says that a list of constraints indexed by the set `model.I` will be created and for each member of `model.I`, the function `ax_constraint_rule` will be called and it will be passed the model object as well as the member of `model.I`.

```
# the next line creates one constraint for each member of the set model.I
model.AxbConstraint = Constraint(model.I, rule=ax_constraint_rule)
```

In the object oriented view of all of this, we would say that `model` object is a class instance of the `AbstractModel` class, and `model.J` is a `Set` object that is contained by this model. Many modeling components in Pyomo can be optionally specified as *indexed components*: collections of components that are referenced using one or more values. In this example, the parameter `model.c` is indexed with set `model.J`.

In order to use this model, data must be given for the values of the parameters. Here is one file that provides data.

```
# one way to input the data in AMPL format
# for indexed parameters, the indexes are given before the value

param m := 1 ;
param n := 2 ;

param a :=
1 1 3
1 2 4
;

param c:=
1 2
2 3
;

param b := 1 1 ;
```

There are multiple formats that can be used to provide data to a Pyomo model, but the AMPL format works well for our purposes because it contains the names of the data elements together with the data. In AMPL data files, text after a pound sign is treated as a comment. Lines generally do not matter, but statements must be terminated with a semi-colon.

For this particular data file, there is one constraint, so the value of `model.m` will be one and there are two variables (i.e., the vector `model.x` is two elements long) so the value of `model.n` will be two. These two assignments are accomplished with standard assignments. Notice that in AMPL format input, the name of the model is omitted.

```
param m := 1 ;
param n := 2 ;
```

There is only one constraint, so only two values are needed for `model.a`. When assigning values to arrays and vectors in AMPL format, one way to do it is to give the index(es) and the the value. The line `1 2 4` causes `model.a[1, 2]` to get the value 4. Since `model.c` has only one index, only one index value is needed so, for example, the line `1 2` causes `model.c[1]` to get the value 2. Line breaks generally do not matter in AMPL format data files, so the assignment of the value for the single index of `model.b` is given on one line since that is easy to read.

```
param a :=
1 1 3
1 2 4
;

param c:=
1 2
2 3
;

param b := 1 1 ;
```

When working with Pyomo (or any other AML), it is convenient to write abstract models in a somewhat more abstract way by using index sets that contain strings rather than index sets that are implied by $1, \dots, m$ or the summation from 1 to n . When this is done, the size of the set is implied by the input, rather than specified directly. Furthermore, the index entries may have no real order. Often, a mixture of integers and indexes and strings as indexes is needed in the same model. To start with an illustration of general indexes, consider a slightly different Pyomo implementation of the model we just presented.

```
# abstract2.py

from __future__ import division
from coopr.pyomo import *

model = AbstractModel()

model.I = Set()
```

```

model.J = Set()

model.a = Param(model.I, model.J)
model.b = Param(model.I)
model.c = Param(model.J)

# the next line declares a variable indexed by the set J
model.x = Var(model.J, domain=NonNegativeReals)

def obj_expression(model):
    return summation(model.c, model.x)

model.OBJ = Objective(rule=obj_expression)

def ax_constraint_rule(model, i):
    # return the expression for the constraint for i
    return sum(model.a[i,j] * model.x[j] for j in model.J) >= model.b[i]

# the next line creates one constraint for each member of the set model.I
model.AxbConstraint = Constraint(model.I, rule=ax_constraint_rule)

```

To get the same instantiated model, the following data file can be used.

```

# abstract2a.dat AMPL format

set I := 1 ;
set J := 1 2 ;

param a :=
1 1 3
1 2 4
;

param c:=
1 2
2 3
;

param b := 1 1 ;

```

However, this model can also be fed different data for problems of the same general form using meaningful indexes.

```

# abstract2.dat AMPL data format

set I := TV Film ;
set J := Graham John Carol ;

param a :=
TV Graham 3
TV John 4.4
TV Carol 4.9
Film Graham 1
Film John 2.4
Film Carol 1.1
;

param c := [*]
Graham 2.2
John 3.1416
Carol 3
;

```

```
param b := TV 1 Film 1 ;
```

2.5 A Simple Concrete Pyomo Model

It is possible to get nearly the same flexible behavior from models declared to be abstract and models declared to be concrete in Pyomo; however, we will focus on a straightforward concrete example here where the data is hard-wired into the model file. Python programmers will quickly realize that the data could have come from other sources.

We repeat the concrete model already given:

$$\begin{array}{ll} \min & 2x_1 + 3x_2 \\ \text{s.t.} & 3x_1 + 4x_2 \geq 1 \\ & x_1, x_2 \geq 0 \end{array}$$

This is implemented as a concrete model as follows:

```
from __future__ import division
from coopr.pyomo import *

model = ConcreteModel()

model.x = Var([1,2], domain=NonNegativeReals)

model.OBJ = Objective(expr = 2*model.x[1] + 3*model.x[2])

model.Constraint1 = Constraint(expr = 3*model.x[1] + 4*model.x[2] >= 1)
```

Although rule functions can also be used to specify constraints and objectives, in this example we use the `expr` option that is available only in concrete models. This option gives a direct specification of the expression.

2.6 Solving the Simple Examples

Pyomo supports modeling and scripting but does not install a solver automatically. In order to solve a model, there must be a solver installed on the computer to be used. If there is a solver, then the `pyomo` command can be used to solve a problem instance.

Suppose that the solver named `glpk` (also known as `glpsol`) is installed on the computer. Suppose further that an abstract model is in the file named `abstract1.py` and a data file for it is in the file named `abstract1.dat`. From the command prompt, with both files in the current directory, a solution can be obtained with the command:

```
pyomo abstract1.py abstract1.dat --solver=glpk
```

Since `glpk` is the default solver, there really is no need specify it so the `--solver` option can be dropped.

Note

There are two dashes before the command line option names such as `solver`.

To continue the example, if `CPLEX` is installed then it can be listed as the solver. The command to solve with `CPLEX` is

```
pyomo abstract1.py abstract1.dat --solver=cplex
```

This yields the following output on the screen:

```
[ 0.00] Setting up Pyomo environment
[ 0.00] Applying Pyomo preprocessing actions
[ 0.07] Creating model
[ 0.15] Applying solver
[ 0.37] Processing results
Number of solutions: 1
Solution Information
  Gap: 0.0
  Status: optimal
  Function Value: 0.6666666666667
Solver results file: results.json
[ 0.39] Applying Pyomo postprocessing actions
[ 0.39] Pyomo Finished
```

The numbers in square brackets indicate how much time was required for each step. Results are written to the file named `results.json`, which has a special structure that makes it useful for post-processing. To see a summary of results written to the screen, use the `--summary` option:

```
pyomo abstract1.py abstract1.dat --solver=cplex --summary
```

To see a list of Pyomo command line options, use:

```
pyomo --help
```

Note

There are two dashes before `help`.

For a concrete model, no data file is specified on the Pyomo command line.

Chapter 3

Sets

3.1 Declaration

Sets can be declared using the `Set` and `RangeSet` functions or by assigning set expressions. The simplest set declaration creates a set and postpones creation of its members:

```
model.A = Set()
```

The `Set` function takes optional arguments such as:

- `doc` = String describing the set
- `dimen` = Dimension of the members of the set
- `filter` = A boolean function used during construction to indicate if a potential new member should be assigned to the set
- `initialize` = A function that returns the members to initialize the set. `ordered` = A boolean indicator that the set is ordered; the default is `False`
- `validate` = A boolean function that validates new member data
- `virtual` = A boolean indicator that the set will never have elements; it is unusual for a modeler to create a virtual set; they are typically used as domains for sets, parameters and variables
- `within` = `Set` used for validation; it is a super-set of the set being declared.

One way to create a set whose members will be two dimensional is to use the `dimen` argument:

```
model.B = Set(dimen=2)
```

To create a set of all the numbers in set `model.A` doubled, one could use

```
def doubleA_init(model):  
    return (i*2 for i in model.A)  
model.C = Set(initialize=DoubleA_init)
```

As an aside we note that as always in Python, there are lot of ways to accomplish the same thing. Also, note that this will generate an error if `model.A` contains elements for which multiplication times two is not defined.

The `initialize` option can refer to a Python set, which can be returned by a function or given directly as in

```
model.D = Set(initialize=['red', 'green', 'blue'])
```

The `initialize` option can also specify a function that is applied sequentially to generate set members. Consider the case of a simple set. In this case, the initialization function accepts a set element number and model and returns the set element associated with that number:

```
def Z_init(model, i):
    if i > 10:
        return Set.End
    return 2*i+1
model.Z = Set(initialize=Z_init)
```

The `Set.End` return value terminates input to the set. Additional information about iterators for set initialization is in [\[Pyomo-Book\]](#).

Note

Data specified in an input file will override the data specified by the initialize options.

If sets are given as arguments to `Set` without keywords, they are interpreted as indexes for an array of sets. For example, to create an array of sets that is indexed by the members of the set `model.A`, use

```
model.E = Set(model.A)
```

Arguments can be combined. For example, to create an array of sets with three dimensional members indexed by set `model.A`, use

```
model.F = Set(model.A, dimen=3)
```

The `initialize` option can be used to create a set that contains a sequence of numbers, but the `RangeSet` function provides a concise mechanism for simple sequences. This function takes as its arguments a start value, a final value, and a step size. If the `RangeSet` has only a single argument, then that value defines the final value in the sequence; the first value and step size default to one. If two values given, they are the first and last value in the sequence and the step size defaults to one. For example, the following declaration creates a set with the numbers 1.5, 5 and 8.5:

```
model.G = RangeSet(1.5, 10, 3.5)
```

3.2 Operations

Sets may also be created by assigning other Pyomo sets as in these examples that also illustrate the set operators union, intersection, difference, and exclusive-or:

```
model.H = model.A
model.I = model.A | model.D # union
model.J = model.A & model.D # intersection
model.K = model.A - model.D # difference
model.L = model.A ^ model.D # exclusive-or
```

The cross-product operator is the asterisk (*). For example, to assign a set the cross product of two other sets, one could use

```
model.K = model.B * model.C
```

or to indicate the the members of a set are restricted to be in the cross product of two other sets, one could use

```
model.K = Set(within=model.B * model.C)
```

The cross-product operator is the asterisk (*). For example, to create a set that contains the cross-product of sets A and B, use

```
model.C = Set(model.A * model.B)
```

to instead create a set that can contain a subset of the members of this cross-product, use

```
model.C = Set(within=model.A * model.B)
```

3.3 Predefined Virtual Sets

For use in specifying domains for sets, parameters and variables, Pyomo provides the following pre-defined virtual sets:

- Any: all possible values
- Reals : floating point values
- PositiveReals: strictly positive floating point values
- NonPositiveReals: non-positive floating point values
- NegativeReals: strictly negative floating point values
- NonNegativeReals: non-negative floating point values
- PercentFraction: floating point values in the interval [0,1]
- Integers: integer values
- PositiveIntegers: positive integer values
- NonPositiveIntegers: non-positive integer values
- NegativeIntegers: negative integer values
- NonNegativeIntegers: non-negative integer values
- Boolean: boolean values, which can be represented as False/True, 0/1, 'False'/'True' and 'F'/'T'
- Binary: same as boolean

For example, if the set `model.M` is declared to be within the virtual set `NegativeIntegers` then an attempt to add anything other than a negative integer will result in an error. Here is the declaration:

```
model.M = Set(within=NegativeIntegers)
```

3.4 Sparse Index Sets

Sets provide indexes for parameters, variables and other sets. Index set issues are important for modelers in part because of efficiency considerations, but primarily because the right choice of index sets can result in very natural formulations that are conducive to understanding and maintenance. Pyomo leverages Python to provide a rich collection of options for index set creation and use.

The choice of how to represent indexes often depends on the application and the nature of the instance data that are expected. To illustrate some of the options and issues, we will consider problems involving networks. In many network applications, it is useful to declare a set of nodes, such as

```
model.Nodes = Set()
```

and then a set of arcs can be created with reference to the nodes.

Consider the following simple version of minimum cost flow problem:

$$\begin{aligned}
 &\text{minimize} && \sum_{a \in \mathcal{A}} c_a x_a \\
 &\text{subject to:} && S_n + \sum_{(i,n) \in \mathcal{A}} x_{(i,n)} \\
 & && -D_n - \sum_{(n,j) \in \mathcal{A}} x_{(n,j)} \quad n \in \mathcal{N} \\
 & && x_a \geq 0, \quad a \in \mathcal{A}
 \end{aligned}$$

where

- Set: Nodes $\equiv \mathcal{N}$
- Set: Arcs $\equiv \mathcal{A} \subseteq \mathcal{N} \times \mathcal{N}$
- Var: Flow on arc (i, j) : $\equiv x_{i,j}, (i, j) \in \mathcal{A}$
- Param: Flow Cost on arc (i, j) : $\equiv c_{i,j}, (i, j) \in \mathcal{A}$
- Param: Demand at node i : $\equiv D_i, i \in \mathcal{N}$
- Param: Supply at node i : $\equiv S_i, i \in \mathcal{N}$

In the simplest case, the arcs can just be the cross product of the nodes, which is accomplished by the definition

```
model.Arcs = model.Nodes * model.Nodes
```

that creates a set with two dimensional members. For applications where all nodes are always connected to all other nodes this may suffice. However, issues can arise when the network is not fully dense. For example, the burden of avoiding flow on arcs that do not exist falls on the data file where high-enough costs must be provided for those arcs. Such a scheme is not very elegant or robust.

For many network flow applications, it might be better to declare the arcs using

```
model.Arcs = Set(within=model.Nodes*model.Nodes)
```

or

```
model.Arcs = Set(dimen=2)
```

where the difference is that the first version will provide error checking as data is assigned to the set elements. This would enable specification of a sparse network in a natural way. But this results in a need to change the `FlowBalance` constraint because as it was written in the simple example, it sums over the entire set of nodes for each node. One way to remedy this is to sum only over the members of the set `model.arcs` as in

```
def FlowBalance_rule(model, node):
    return model.Supply[node] \
        + sum(model.Flow[i, node] for i in model.Nodes if (i,node) in model.Arcs) \
        - model.Demand[node] \
        - sum(model.Flow[node, j] for j in model.Nodes if (j,node) in model.Arcs) \
        == 0
```

This will be OK unless the number of nodes becomes very large for a sparse network, then the time to generate this constraint might become an issue (admittely, only for very large networks, but such networks do exist).

Another method, which comes in handy in many network applications, is to have a set for each node that contain the nodes at the other end of arcs going to the node at hand and another set giving the nodes on out-going arcs. If these sets are called `model.NodesIn` and `model.NodesOut` respectively, then the flow balance rule can be re-written as

```
def FlowBalance_rule(model, node):
    return model.Supply[node] \
        + sum(model.Flow[i, node] for i in model.NodesIn[node]) \
        - model.Demand[node] \
        - sum(model.Flow[node, j] for j in model.NodesOut[node]) \
        == 0
```

The data for `NodesIn` and `NodesOut` could be added to the input file, and this may be the most efficient option.

For all but the largest networks, rather than reading `Arcs`, `NodesIn` and `NodesOut` from a data file, it might be more elegant to read only `Arcs` from a data file and declare `model.NodesIn` with an `initialize` option specifying the creation as follows:

```
def NodesIn_init(model, node):
    retval = []
    for (i,j) in model.Arcs:
        if j == node:
            retval.append(i)
    return retval
model.NodesIn = Set(model.Nodes, initialize=NodesIn_init)
```

with a similar definition for `model.NodesOut`. This code creates a list of sets for `NodesIn`, one set of nodes for each node. The full model is:

```
# Isinglecomm.py
# NodesIn and NodesOut are intialized using the Arcs
from coopr.pyomo import *

model = AbstractModel()

model.Nodes = Set()
model.Arcs = Set(dimen=2)

def NodesOut_init(model, node):
    retval = []
    for (i,j) in model.Arcs:
        if i == node:
            retval.append(j)
    return retval
model.NodesOut = Set(model.Nodes, initialize=NodesOut_init)

def NodesIn_init(model, node):
    retval = []
    for (i,j) in model.Arcs:
        if j == node:
            retval.append(i)
    return retval
model.NodesIn = Set(model.Nodes, initialize=NodesIn_init)

model.Flow = Var(model.Arcs, domain=NonNegativeReals)
model.FlowCost = Param(model.Arcs)

model.Demand = Param(model.Nodes)
model.Supply = Param(model.Nodes)

def Obj_rule(model):
    return summation(model.FlowCost, model.Flow)
model.Obj = Objective(rule=Obj_rule, sense=minimize)

def FlowBalance_rule(model, node):
    return model.Supply[node] \
        + sum(model.Flow[i, node] for i in model.NodesIn[node]) \
        - model.Demand[node] \
        - sum(model.Flow[node, j] for j in model.NodesOut[node]) \
        == 0
model.FlowBalance = Constraint(model.Nodes, rule=FlowBalance_rule)
```

for this model, a toy data file would be:

```
# Isinglecomm.dat: data for Isinglecomm.py

set Nodes := CityA CityB CityC ;

set Arcs :=
```

```

CityA CityB
CityA CityC
CityC CityB
;

param : FlowCost :=
CityA CityB 1.4
CityA CityC 2.7
CityC CityB 1.6
;

param Demand :=
CityA 0
CityB 1
CityC 1
;

param Supply :=
CityA 2
CityB 0
CityC 0
;

```

3.4.1 Sparse Index Sets Example

One may want to have a constraint that holds

```
for i in model.I, k in model.K, v in model.V[k]
```

There are many ways to accomplish this, but one good way is to create a set of tuples composed of all of `model.k`, `model.V[k]` pairs. This can be done as follows:

```
def kv_init(model):
    return ((k,v) for k in model.K for v in model.V[k])
model.KV=Set(dimen=2, initialize=kv_init)
```

So then if there was a constraint defining rule such as

```
def MyC_rule(model, i, k, v):
    return ...
```

Then a constraint could be declared using

```
model.MyConstraint = Constraint(model.I,model.KV,rule=c1Rule)
```

Here is the first few lines of a model that illustrates this:

```

from coopr.pyomo import *

model = AbstractModel()

model.I=Set()
model.K=Set()
model.V=Set(model.K)

def kv_init(model):
    return ((k,v) for k in model.K for v in model.V[k])
model.KV=Set(dimen=2, initialize=kv_init)

model.a = Param(model.I, model.K)

```

```
model.y = Var(model.I)
model.x = Var(model.I, model.KV)

#include a constraint
#x[i,k,v] <= a[i,k]*y[i], for i in model.I, k in model.K, v in model.V[k]

def c1Rule(model,i,k,v):
    return model.x[i,k,v] <= model.a[i,k]*model.y[i]
model.c1 = Constraint(model.I,model.KV,rule=c1Rule)
```

Chapter 4

Parameters

The word "parameters" is used in many settings. When discussing a Pyomo model, we use the word to refer to data that must be provided in order to find an optimal (or good) assignment of values to the decision variables. Parameters are declared with the `Param` function, which takes arguments that are very similar to the `Set` function. For example, the following code snippet declares sets `model.A`, `model.B` and then a parameter array `model.P` that is indexed by `model.A`:

```
model.A = Set()
model.B = Set()
model.P = Param(model.A, model.B)
```

In addition to sets that serve as indexes, the `Param` function takes the following command options:

- `default` = The value absent any other specification.
- `doc` = String describing the parameter
- `initialize` = A function (or Python object) that returns the members to initialize the parameter values.
- `rule` = (this is a synonym for `initilize`)
- `validate` = A boolean function with arguments that are the prospective parameter value, the parameter indices and the model.
- `within` = `Set` used for validation; it specifies the domain of the parameter values.

These options perform in the same way as they do for `Set`. For example, suppose that `Model.A = RangeSet(1, 3)`, then there are many ways to create a parameter that is a square matrix with 9, 16, 25 on the main diagonal zeros elsewhere, here are two ways to do it. First using a Python object to initialize:

```
v={}
v[1,1] = 9
v[2,2] = 16
v[3,3] = 25
model.S = Param(model.A, model.A, initialize=v, default=0)
```

And now using an initialization rule that is automatically called once for each index tuple (remember that we are assuming that `model.A` contains 1,2,3)

```
def s_init(model, i, j):
    if i == j:
        return i*i
    else:
        return 0.0
model.S = Param(model.A, model.A, rule=s_init)
```

In this example, the index set contained integers, but index sets need not be numeric. It is very common to use strings.

Note

Data specified in an input file will override the data specified by the initialize options.

Chapter 5

Variables

Variables are intended to ultimately be given values by an optimization package. They are declared and optionally bounded, given initial values, and documented using the Pyomo `Var` function. If index sets are given as arguments to this function they are used to index the variable, other optional directives include:

- `bounds` = A function (or Python object) that gives a (lower,upper) bound pair for the variable
- `domain` = A set that is a super-set of the values the variable can take on.
- `initialize` = A function (or Python object) that gives a starting value for the variable; this is particularly important for non-linear models
- `within` = (synonym for `domain`)

The following code snippet illustrates some aspects of these options by declaring a *singleton* (i.e. unindexed) variable named `model.LumberJack` that will take on real values between zero and 6 and it initialized to be 1.5:

```
model.LumberJack = Var(within=NonNegativeReals, bounds=(0,6), initialize=1.5)
```

Instead of the `initialize` option, initialization is sometimes done with a Python assignment statement as in

```
model.LumberJack = 1.5
```

For indexed variables, bounds and initial values are often specified by a rule (a Python function) that itself may make reference to parameters or other data. The formal arguments to these rules begins with the model followed by the indexes. This is illustrated in the following code snippet that makes use of Python dictionaries declared as `lb` and `ub` that are used by a function to provide bounds:

```
model.A = Set(initialize=['Scones', 'Tea'])
lb = {'Scones':2, 'Tea':4}
ub = {'Scones':5, 'Tea':7}
def fb(model, i):
    return (lb[i], ub[i])
model.PriceToCharge = Var(model.A, domain=PositiveInteger, bounds=fb)
```

Note

Many of the pre-defined virtual sets that are used as domains imply bounds. A strong example is the set `Boolean` that implies bounds of zero and one.

Chapter 6

Objectives

An objective is a function of variables that returns a value that an optimization package attempts to maximize or minimize. The `Objective` function in Pyomo declares an objective. Although other mechanisms are possible, this function is typically passed the name of another function that gives the expression. Here is a very simple version of such a function that assumes `model.x` has previously been declared as a `Var`:

```
def ObjRule(model):
    return 2*model.x[1] + 3*model.x[2]
model.g = Objective(rule=ObjRule)
```

It is more common for an objective function to refer to parameters as in this example that assumes that `model.p` has been declared as a parameters and that `model.x` has been declared with the same index set, while `model.y` has been declared as a singleton:

```
def profrul(model):
    return summation(model.p, model.x) + model.y
model.Obj = Objective(rule=ObjRule, sense=maximize)
```

This example uses the `sense` option to specify maximization. The default sense is `minimize`.

Chapter 7

Constraints

Most constraints are specified using equality or inequality expressions that are created using a rule, which is a Python function. For example, if the variable `model.x` has the indexes *butter* and *scones*, then this constraint limits the sum for them to be exactly three:

```
def teaOKrule(model):
    return(model.x['butter'] + model.x['scones'] == 3)
model.TeaConst = Constraint(rule=teaOKrule)
```

Instead of expressions involving equality (`==`) or inequalities (`<=` or `>=`), constraints can also be expressed using a 3-tuple if the form `(lb, expr, ub)` where `lb` and `ub` can be `None`, which is interpreted as `lb <= expr <= ub`. Variables can appear only in the middle `expr`. For example, the following two constraint declarations have the same meaning:

```
model.x = Var()

def aRule(model):
    return model.x >= 2
Boundx = Constraint(rule=aRule)

def bRule(model):
    return (2, model.x, None)
Boundx = Constraint(rule=bRule)
```

For this simple example, it would also be possible to declare `model.x` with a `bound` option to accomplish the same thing.

Constraints (and objectives) can be indexed by lists or sets. When the declaration contains lists or sets as arguments, the elements are iteratively passed to the rule function. If there is more than one, then the cross product is sent. For example the following constraint could be interpreted as placing a budget of i on the i^{th} item to buy where the cost per item is given by the parameter `model.a`:

```
model.A = RangeSet(1,10)
model.a = Param(model.A, within=PositiveReals)
model.ToBuy = Var(model.A)
def bud_rule(model, i):
    return model.a[i]*model.ToBuy[i] <= i
aBudget = Constraint(model.A)
```

Note

Python and Pyomo are case sensitive so `model.a` is not the same as `model.A`.

Chapter 8

Expressions

8.1 Rules to Generate Expressions

Both objectives and constraints make use of rules to generate expressions. These are Python functions that return the appropriate expression. These are first-class functions that can access global data as well as data passed in, including the model object.

Operations on model elements results in expressions, which seems natural in expression like the constraints we have seen so far. It is also possible to build up expressions. The following example illustrates this along with a reference to global Python data in the form of a Python variable called `switch`:

```
switch = 3

model.A = RangeSet(1, 10)
model.c = Param(model.A)
model.d = Param()
model.x = Var(model.A, domain=Boolean)

def pi_rule(model)
    accexpr = summation(model.c, model.x)
    if switch >= 2:
        accexpr = accexpr - model.d
    return accexpr >= 0.5
PieSlice = Constraint(rule=pi_rule)
```

In this example, the constraint that is generated depends on the value of the Python variable called `switch`. If the value is 2 or greater, then the constraint is `summation(model.c, model.x) - model.d >= 0.5`; otherwise, the `model.d` term is not present.



Caution

Because model elements result in expressions, not values, the following does not work as expected in an abstract model!

```
model.A = RangeSet(1, 10)
model.c = Param(model.A)
model.d = Param()
model.x = Var(model.A, domain=Boolean)

def pi_rule(model)
    accexpr = summation(model.c, model.x)
    if model.d >= 2: # NOT in an abstract model!!
        accexpr = accexpr - model.d
```

```

    return accexpr >= 0.5
PieSlice = Constraint(rule=pi_rule)

```

The trouble is that `model.d >= 2` results in an expression, not its evaluated value. Instead use `if value(model.d) >= 2`

8.2 Piecewise Linear Expressions

Pyomo has facilities to add piecewise constraints of the form $y=f(x)$ for a variety of forms of the function f .

The piecewise types other than `SOS2`, `BIGM_SOS1`, `BIGM_BIN` are implemented as described in the paper [\[Vielma_et_al\]](#).

There are two basic forms for the declaration of the constraint:

```

model.pwconst = Piecewise(index_1, ..., index_n, yvar, xvar, **Keywords)
model.pwconst = Piecewise(yvar, xvar, **Keywords)

```

where `pwconst` can be replaced by a name appropriate for the application. The choice depends on whether the x and y variables are indexed. If so, they must have the same index sets and these sets are given as the first arguments.

KEYWORDS:

- `pw_pts={},[],()` A dictionary of lists (keys are index set) or a single list (for the non-indexed case or when an identical set of breakpoints is used across all indices) defining the set of domain breakpoints for the piecewise linear function. NOTE: `pw_pts` is always required. These give the breakpoints for the piecewise function and are expected to full span the bounds for the independent variable(s).
- `pw_repn=<Option>` Indicates the type of piecewise representation to use. This can have a major impact on solver performance. Options: (Default 'SOS2')
 - 'SOS2' - Standard representation using `sos2` constraints.
 - 'BIGM_BIN' - BigM constraints with binary variables. The theoretically tightest M values are automatically determined.
 - 'BIGM_SOS1' - BigM constraints with `sos1` variables. The theoretically tightest M values are automatically determined.
 - 'DCC' - Disaggregated convex combination model.
 - 'DLOG' - Logarithmic disaggregated convex combination model.
 - 'CC' - Convex combination model.
 - 'LOG' - Logarithmic branching convex combination.
 - 'MC' - Multiple choice model.
 - 'INC' - Incremental (delta) method. NOTE: Step functions are supported for all but the two BIGM options. Refer to the `force_pw` option.
- `pw_constr_type= <Option>` Indicates the bound type of the piecewise function. Options:
 - 'UB' - y variable is bounded above by piecewise function
 - 'LB' - y variable is bounded below by piecewise function
 - 'EQ' - y variable is equal to the piecewise function
- `f_rule=f(model,i,j,...,x), {}, [], ()`
 An object that returns a numeric value that is the range value corresponding to each piecewise domain point. For functions, the first argument must be a Pyomo model. The last argument is the domain value at which the function evaluates (Not a Pyomo `Var`). Intermediate arguments are the corresponding indices of the Piecewise component (if any). Otherwise, the object can be a dictionary of lists/tuples (with keys the same as the indexing set) or a single list/tuple (when no indexing set is used or when all indices use an identical piecewise function). Examples:

```

# A function that changes with index
def f(model, j, x):
    if (j == 2):
        return x**2 + 1.0
    else:
        return x**2 + 5.0

# A nonlinear function
f = lambda model, x: return exp(x) + value(model.p)
    (model.p is a Pyomo Param)

# A step function
f = [0, 0, 1, 1, 2, 2]

```

- `force_pw=True/False`
Using the given function rule and `pw_pts`, a check for convexity/concavity is implemented. If (1) the function is convex and the piecewise constraints are lower bounds or if (2) the function is concave and the piecewise constraints are upper bounds then the piecewise constraints will be substituted for linear constraints. Setting `force_pw=True` will force the use of the original piecewise constraints even when one of these two cases applies.
- `warning_tol=<float>`
To aid in debugging, a warning is printed when consecutive slopes of piecewise segments are within `<warning_tol>` of each other. Default=`1e-8`
- `warn_domain_coverage=True/False`
Print a warning when the feasible region of the domain variable is not completely covered by the piecewise breakpoints. Default=`True`
- `unbounded_domain_var=True/False`
Allow an unbounded or partially bounded Pyomo Var to be used as the domain variable. Default=`False` NOTE: This does not imply unbounded piecewise segments will be constructed. The outermost piecewise breakpoints will bound the domain variable at each index. However, the Var attributes `.lb` and `.ub` will not be modified.

Here is an example of an assignment to a Python dictionary variable that has keywords for a piecewise constraint:

```
kwds = {'pw_constr_type': 'EQ', 'pw_repn': 'SOS2', 'sense': maximize, 'force_pw': True}
```

Here is a simple example based on the [abstract2.py](#) example given early. In this new example, the objective function is the sum of `c` times `x` to the fourth. In this example, the keywords are passed directly to the `Piecewise` function without being assigned to a dictionary variable. The upper bound on the `x` variables was chosen whimsically just to make the example. The important thing to note is that variables that are going to appear as the independent variable in a piecewise constraint must have bounds.

```

# abstract2piece.py
# Similar to abstract2.py, but the objective is now c times x to the fourth power

from __future__ import division
from coopr.pyomo import *

model = AbstractModel()

model.I = Set()
model.J = Set()

model.a = Param(model.I, model.J)
model.b = Param(model.I)
model.c = Param(model.J)

model.Topx = Param(default=6.1) # range of x variables

```

```
# the next line declares a variable indexed by the set J
model.x = Var(model.J, domain=NonNegativeReals, bounds=(0,model.Topx))
model.y = Var(model.J, domain=NonNegativeReals)

# to avoid warnings, we set breakpoints beyond the bounds
bpts = []
for i in range(0,int(model.Topx+1)): # simply the integer lattice
    bpts.append(i)

def f4(model, j, xp):
    # we not need j, but it is passed as the index for the constraint
    return xp**4

model.ComputeObj = Piecewise(model.J, model.y, model.x, pw_pts=bpts, f_rule=f4, ←
    pw_constr_type='EQ')

def obj_expression(model):
    return summation(model.c, model.y)

model.OBJ = Objective(rule=obj_expression)

def ax_constraint_rule(model, i):
    # return the expression for the constraint for i
    return sum(model.a[i,j] * model.x[j] for j in model.J) >= model.b[i]

# the next line creates one constraint for each member of the set model.I
model.AxbConstraint = Constraint(model.I, rule=ax_constraint_rule)
```

Chapter 9

Data Input

For abstract models, Pyomo supports data input from a data command file or using a `DataPortal` object.

9.1 Data Command Files

The following commands can be used in data command files:

- `set` declares set data.
- `param` declares a table of parameter data, which can also include the declaration of the set data used to index parameter data.
- `import` loads data from an external data source such as ASCII table files, CSV files, ranges in spreadsheets, and database tables.
- `include` specifies a data command file that is to be processed immediately.
- The `data` and `end` commands do not perform any actions, but they provide compatibility with AMPL scripts that define data commands.
- `namespace` defines groupings of data commands.

Note

The syntax of the `set` and `param` data commands are adapted from AMPL's data commands. However, other Pyomo data commands do not directly correspond to AMPL data commands. In particular, the syntax of the AMPL `table` command allows the user to specify complex mappings from table data values to corresponding model parameters and sets. The corresponding Pyomo `import` command supports much simpler mappings. Complex mappings are accomplished in Pyomo via scripting.

The `namespace` command allows data commands to be organized into named groups that can be enabled from the `pyomo` command line. For example, consider again the [abstract2.py](#) example. Suppose that the cost data shown in [abstract2.dat](#) were valid only under certain circumstances that we will label as "TerryG" and that there would be different cost data under circumstances that we will label "JohnD." This could be represented using the following data file:

```
# abs2namespace.dat AMPL format with namespaces

set I := TV Film ;
set J := Graham John Carol ;

param a :=
TV   Graham 3
TV   John   4.4
TV   Carol  4.9
```

```

Film Graham 1
Film John 2.4
Film Carol 1.1
;

namespace TerryG {
  param c := [*]
  Graham 2.2
  John 3.1416
  Carol 3
;
}

namespace JohnD {
  param c := [*]
  Graham 2.7
  John 3
  Carol 2.1
;
}

param b := TV 1 Film 1 ;

```

To use this data file with `abstract2.py`, a namespace must be indicated on the command line. To select the "TerryG" data specification, `--namespace TerryG` would be added to the command line. For example:

```
pyomo abstract2.py abs2nspace.dat --namespace TerryG --solver=cplex
```

If the `--namespace` option is omitted, then no data will be given for `model.c` (and no default was given for `model.c`). In other words, there is no default namespace selection.

The option `-ns` (with one dash) is an alias for `--namespace` (which needs two dashes) Multiple namespaces can be selected by giving multiple `--namespace` or `-ns` arguments on the Pyomo command line.

9.2 DataPortal Objects

The `import` and `export` Pyomo data commands can be used to load set and table data from a variety of data sources. Pyomo's `DataPortal` object provides this same functionality for users who work with Python scripts. A `DataPortal` object manages the process of loading data from different data sources, and it is used to construct model instances in a standard manner. Similarly, a `DataPortal` object can be used to store model data externally in a standard manner.

Note

Pyomo also supports the `ModelData` object, which provides a narrow set of capabilities than is supported by the `DataPortal` object. The use of `ModelData` objects is deprecated.

Note

The Pyomo data commands `import` and `export` correspond to the `DataPortal` methods `load` and `store`. This discrepancy is due to the fact that `import` is a reserved word in Python. Thus, we cannot define an `import` method in the `DataPortal` class. (TODO: we need to create `load` and `store` data commands, and deprecated the `import` and `export` data commands.)

9.2.1 Loading Data

The `load` method can be used to load data into Pyomo models from a variety of sources and formats. The most common format is a table representation of set and parameter data. For example, consider the file `A.tab`, which defines a simple set:

```
A
A1
A2
A3
```

The following example illustrates how a `DataPortal` object can be used to load this data into a model:

```
model = AbstractModel()
model.A = Set()

data = DataPortal()
data.load(filename='tab/A.tab', set=model.A)
instance = model.create(data)
```

The `load` method opens the data file, processes it, and loads the data in a format that is then used to construct a model instance. The `load` method can be called multiple times to load data for different sets or parameters, or to override data processed earlier.

Note

Subsequent examples omit the model declaration and instance creation.

In the previous example, the `set` option is used to define the model component that is loaded with the set data. If the data source defines a table of data, then this option is used to specify data for a multi-dimensional set. For example, consider the file `D.tab`:

```
A B
A1 1
A1 2
A1 3
A2 1
A2 2
A2 3
A3 1
A3 2
A3 3
```

If a two-dimensional set is declared, then it can be loaded with the same syntax:

```
model.A = Set(dimen=2)

data.load(filename='tab/C.tab', set=model.A)
```

This example also illustrates that the column titles do not directly impact the process of loading data. Column titles are only used to select columns that are included in the table that is loaded (see below).

The `param` option is used to define the a parameter component that is loaded with data. The simplest parameter is a singleton. For example, consider the file `Z.tab`:

```
1.1
```

This data is loaded with the following syntax:

```
model.z = Param()

data.load(filename='tab/Z.tab', param=model.z)
```

Indexed parameters can be defined from table data. For example, consider the file `Y.tab`:

```
A Y
A1 3.3
A2 3.4
A3 3.5
```

The parameter `y` is loaded with the following syntax:

```
model.A = Set(initialize=['A1', 'A2', 'A3', 'A4'])
model.y = Param(model.A)

data.load(filename='tab/Y.tab', param=model.y)
```

Pyomo assumes that the parameter values are defined on the rightmost column; the column names are not used to specify the index and parameter data (see below). In this file, the `A` column contains the index values, and the `Y` column contains the parameter values.

Similarly, multiple parameters can be initialized at once by specifying a list (or tuple) of component parameters. For example, consider the file `XW.tab`:

```
A X W
A1 3.3 4.3
A2 3.4 4.4
A3 3.5 4.5
```

The parameters `x` and `w` are loaded with the following syntax:

```
model.A = Set(initialize=['A1', 'A2', 'A3', 'A4'])
model.x = Param(model.A)
model.w = Param(model.A)

data.load(filename='tab/XW.tab', param=(model.x, model.w))
```

Note that the data for set `A` is predefined in this example. The index set can be loaded with the parameter data using the `index` option:

```
model.A = Set()
model.x = Param(model.A)
model.w = Param(model.A)

data.load(filename='tab/XW.tab', param=(model.x, model.w), index=model.A)
```

We have previously noted that the column names are not used to define the set and parameter data. The `select` option is used to define the columns in the table that are used to load data. This option specifies a list (or tuple) of column names that are used, in that order, to form the table that defines the component data.

For example, consider the following load declaration:

```
model.A = Set()
model.w = Param(model.A)

data.load(filename='tab/XW.tab', select=('A', 'W'), param=model.w, index=model.A)
```

The columns `A` and `W` are selected from the file `XW.tab`, and a single parameter is defined.

Note

The `load` method allows for a variety of other options that are supported by the `add` method for `ModelData` objects. See the [\[PyomoBook\]](#) for a detailed description of these options.

Chapter 10

The `pyomo` Command

The `pyomo` command is issued to the DOS prompt or a Unix shell. To see a list of Pyomo command line options, use:

```
pyomo --help
```

Note

There are two dashes before `help`.

In this section we will detail some of the options.

10.1 Passing Options to a Solver

To pass arguments to a solver, use the Pyomo argument `--solver-option=` followed by an argument that is a string to be sent to the solver (perhaps with dashes added by Coopr). So for most MIP solvers, the mip gap can be set using

```
--ef-solver-options= "mipgap=0.01 "
```

Multiple options are separated by a space. For example, to specify that the solver is GLPK, then to specify a mipgap of two percent and the GLPK cuts option, use

```
--solver=glpk --solver-options="mipgap=0.02 cuts"
```

If there are multiple "levels" to the keyword, as is the case for some Gurobi and CPLEX options, the tokens are separated by underscore. For example, `mip cuts all` would be specified as `mip_cuts_all`. For another example, to set the solver to be CPLEX, then to set a mip gap of one percent and to specify `y` for the sub-option `numerical` to the option `emphasis` use

```
--solver=cplex --solver-options="mipgap=0.001 emphasis_numerical=y"
```

10.2 Troubleshooting

Many of things that can go wrong are covered by error messages, but sometimes they can be confusing or do not provide enough information. Depending on what the troubles are, there might be ways to get a little additional information.

If there are syntax errors in the model file, for example, it can occasionally be helpful to get error messages directly from the Python interpreter rather than through Pyomo. Suppose the name of the model file is `scuc.py`, then

```
python scuc.py
```

can sometimes give useful information for fixing syntax errors.

When there are no syntax errors, but there troubles reading the data or generating the information to pass to a solver, then the `--verbose` option provides a trace of the execution of Pyomo. The user should be aware that for some models this option can generate a lot of output.

If there are troubles with solver (i.e., after Pyomo has output "Applying Solver"), it is often helpful to use the option `--stream-solver` that causes the solver output to be displayed rather than trapped. (See [Solver Display](#) for information about getting this output in a script). Advanced users may wish to examine the files that are generated to be passed to a solver. The type of file generated is controlled by the `--solver-io` option and the `--keepfiles` option instructs pyomo to keep the files and output their names.

10.3 Direct Interfaces to Solvers

In many applications, the default solver interface works well. However, in some cases it is useful to specify the interface using the `solver-io` option. For example, if the solver supports a direct Python interface, then the option would be specified on the command line as

```
--solver-io=python
```

Here are some of the choices:

- `lp`: generate a standard linear programming format file with filename extension `lp`
- `nlp`: generate a file with a standard format that supports linear and nonlinear optimization with filename extension `nlp`
- `os`: generate an OSiL format XML file.
- `python`: use the direct Python interface.

Note that not all solvers support all interfaces.

Chapter 11

PySP Overview

This chapter describes PySP: (Pyomo Stochastic Programming), where parameters are allowed to be uncertain.

11.1 Overview of Modeling Components and Processes

The sequence of activities is typically the following:

- Create a deterministic model and declare components
- Develop base-case data for the deterministic model
- Test, verify and validate the deterministic model
- Model the stochastic processes
- Develop a way to generate scenarios (in the form of a tree if there are more than two stages)
- Create the data files need to describe the stochastics
- Use PySP to solve stochastic problem

When viewed from the standpoint of file creation, the process is

- Create an abstract model for the deterministic problem in a file called `ReferenceModel.py`
- Specify data for this model in a file called `ReferenceModel.dat`
- Specify the stochastics in a file called `ScenarioStructure.dat`
- Specify scenario data

11.2 Birge and Louveaux's Farmer Problem

Birge and Louveaux [[BirgeLouveauxBook](#)] make use of the example of a farmer who has 500 acres that can be planted in wheat, corn or sugar beets, at a per acre cost of 150, 230 and 260 (Euros, presumably), respectively. The farmer needs to have at least 200 tons of wheat and 240 tons of corn to use as feed, but if enough is not grown, those crops can be purchased for 238 and 210, respectively. Corn and wheat grown in excess of the feed requirements can be sold for 170 and 150, respectively. A price of 36 per ton is guaranteed for the first 6000 tons grown by any farmer, but beets in excess of that are sold for 10 per ton. The yield is 2.5, 3, and 20 tons per acre for wheat, corn and sugar beets, respectively.

11.2.1 ReferenceModel.py

So far, this is a deterministic problem because we are assuming that we know all the data. The Pyomo model for this problem shown here is in the file `ReferenceModel.py` in the sub-directory `examples/pysp/farmer/models` that is distributed with Coopr.

```
# Farmer: rent out version has a singleton root node var
# note: this will minimize
#
# Imports
#

from __future__ import division
from coopr.pyomo import *

#
# Model
#

model = AbstractModel()

#
# Parameters
#

model.CROPS = Set()

model.TOTAL_ACREAGE = Param(within=PositiveReals)

model.PriceQuota = Param(model.CROPS, within=PositiveReals)

model.SubQuotaSellingPrice = Param(model.CROPS, within=PositiveReals)

def super_quota_selling_price_validate (model, value, i):
    return model.SubQuotaSellingPrice[i] >= model.SuperQuotaSellingPrice[i]

model.SuperQuotaSellingPrice = Param(model.CROPS, validate= ←
    super_quota_selling_price_validate)

model.CattleFeedRequirement = Param(model.CROPS, within=NonNegativeReals)

model.PurchasePrice = Param(model.CROPS, within=PositiveReals)

model.PlantingCostPerAcre = Param(model.CROPS, within=PositiveReals)

model.Yield = Param(model.CROPS, within=NonNegativeReals)

#
# Variables
#

model.DevotedAcreage = Var(model.CROPS, bounds=(0.0, model.TOTAL_ACREAGE))

model.QuantitySubQuotaSold = Var(model.CROPS, bounds=(0.0, None))
model.QuantitySuperQuotaSold = Var(model.CROPS, bounds=(0.0, None))

model.QuantityPurchased = Var(model.CROPS, bounds=(0.0, None))

model.FirstStageCost = Var()
model.SecondStageCost = Var()

#
```

```

# Constraints
#

def ConstrainTotalAcreage_rule(model):
    return summation(model.DevotedAcreage) <= model.TOTAL_ACREAGE

model.ConstrainTotalAcreage = Constraint(rule=ConstrainTotalAcreage_rule)

def EnforceCattleFeedRequirement_rule(model, i):
    return model.CattleFeedRequirement[i] <= (model.Yield[i] * model.DevotedAcreage[i]) + <-
        model.QuantityPurchased[i] - model.QuantitySubQuotaSold[i] - model.<-
        QuantitySuperQuotaSold[i]

model.EnforceCattleFeedRequirement = Constraint(model.CROPS)

def LimitAmountSold_rule(model, i):
    return model.QuantitySubQuotaSold[i] + model.QuantitySuperQuotaSold[i] - (model.Yield[i <-
        ] * model.DevotedAcreage[i]) <= 0.0

model.LimitAmountSold = Constraint(model.CROPS)

def EnforceQuotas_rule(model, i):
    return (0.0, model.QuantitySubQuotaSold[i], model.PriceQuota[i])

model.EnforceQuotas = Constraint(model.CROPS)

#
# Stage-specific cost computations
#

def ComputeFirstStageCost_rule(model):
    return model.FirstStageCost - summation(model.PlantingCostPerAcre, model.DevotedAcreage <-
        ) == 0.0

model.ComputeFirstStageCost = Constraint()

def ComputeSecondStageCost_rule(model):
    expr = summation(model.PurchasePrice, model.QuantityPurchased)
    expr -= summation(model.SubQuotaSellingPrice, model.QuantitySubQuotaSold)
    expr -= summation(model.SuperQuotaSellingPrice, model.QuantitySuperQuotaSold)
    return (model.SecondStageCost - expr) == 0.0

model.ComputeSecondStageCost = Constraint()

#
# Objective
#

def Total_Cost_Objective_rule(model):
    return model.FirstStageCost + model.SecondStageCost

model.Total_Cost_Objective = Objective(sense=minimize)

```

11.2.2 ReferenceModel.dat

The data introduced here are in the file ReferenceModel.dat in the sub-directory examples/pysp/farmer/scenariodata that is distributed with Coopr.

```
set CROPS := WHEAT CORN SUGAR_BEETS ;
```



```

set Scenarios := BelowAverageScenario
               AverageScenario
               AboveAverageScenario ;

param ScenarioLeafNode :=
    BelowAverageScenario BelowAverageNode
    AverageScenario      AverageNode
    AboveAverageScenario AboveAverageNode ;

set StageVariables[FirstStage] := DevotedAcreage[*] ;
set StageVariables[SecondStage] := QuantitySubQuotaSold[*]
                                   QuantitySuperQuotaSold[*]
                                   QuantityPurchased[*] ;

param StageCostVariable := FirstStage FirstStageCost
                           SecondStage SecondStageCost ;

```

This data file is verbose and somewhat redundant, but in most applications it is generated by software rather than by a person, so this is not an issue. Generally, the left-most part of each expression (e.g. “set Stages :=”) is required and uses reserved words (e.g., Stages) and the other names are supplied by the user (e.g., “FirstStage” could be any name). Every assignment is terminated with a semi-colon. We will now consider the assignments in this file one at a time.

The first assignments provides names for the stages and the words "set Stages" are required, as are the := symbols. Any names can be used. In this example, we used "FirstStage" and "SecondStage" but we could have used "EtapPrimero" and "ZweiteEtage" if we had wanted to. Whatever names are given here will continue to be used to refer to the stages in the rest of the file. The order of the names is important. A simple way to think of it is that generally, the names must be in time order (technically, they need to be in order of information discovery, but that is usually time-order). Stages refers to decision stages, which may, or may not, correspond directly with time stages. In the farmer example, decisions about how much to plant are made in the first stage and "decisions" (which are pretty obvious, but which are decision variables nonetheless) about how much to sell at each price and how much needs to be bought are second stage decisions because they are made after the yield is known.

```

set Stages := FirstStage SecondStage ;

```

Node names are constructed next. The words "set Nodes" are required, but any names may be assigned to the nodes. In two stage stochastic problems there is a root node, which we chose to name "RootNode" and then there is a node for each scenario.

```

set Nodes := RootNode
            BelowAverageNode
            AverageNode
            AboveAverageNode ;

```

Nodes are associated with time stages with an assignment beginning with the required words "param Nodestage." The assignments must make use of previously defined node and stage names. Every node must be assigned a stage.

```

param NodeStage := RootNode      FirstStage
                  BelowAverageNode SecondStage
                  AverageNode    SecondStage
                  AboveAverageNode SecondStage ;

```

The structure of the scenario tree is defined using assignment of children to each node that has them. Since this is a two stage problem, only the root node has children. The words "param Children" are required for every node that has children and the name of the node is in square brackets before the colon-equals assignment symbols. A list of children is assigned.

```

set Children[RootNode] := BelowAverageNode
                          AverageNode
                          AboveAverageNode ;

```

The probability for each node, conditional on observing the parent node is given in an assignment that begins with the required words "param ConditionalProbability." The root node always has a conditional probability of 1, but it must always be given anyway. In this example, the second stage nodes are equally likely.

```
param ConditionalProbability := RootNode      1.0
                             BelowAverageNode 0.33333333
                             AverageNode      0.33333334
                             AboveAverageNode 0.33333333 ;
```

Scenario names are given in an assignment that begins with the required words "set Scenarios" and provides a list of the names of the scenarios. Any names may be given. In many applications they are given unimaginative names generated by software such as "Scen1" and the like. In this example, there are three scenarios and the names reflect the relative values of the yields.

```
set Scenarios := BelowAverageScenario
                AverageScenario
                AboveAverageScenario ;
```

Leaf nodes, which are nodes with no children, are associated with scenarios. This assignment must be one-to-one and it is initiated with the words "param ScenarioLeafNode" followed by the colon-equals assignment characters.

```
param ScenarioLeafNode :=
    BelowAverageScenario BelowAverageNode
    AverageScenario      AverageNode
    AboveAverageScenario AboveAverageNode ;
```

Variables are associated with stages using an assignment that begins with the required words "set StageVariables" and the name of a stage in square brackets followed by the colon-equals assignment characters. Variable names that have been defined in the file `ReferenceModel.py` can be assigned to stages. Any variables that are not assigned are assumed to be in the last stage. Variable indexes can be given explicitly and/or wildcards can be used. Note that the variable names appear without the prefix "model." In the farmer example, `DevotedAcreage` is the only first stage variable.

```
set StageVariables[FirstStage] := DevotedAcreage[*] ;
set StageVariables[SecondStage] := QuantitySubQuotaSold[*]
                                   QuantitySuperQuotaSold[*]
                                   QuantityPurchased[*] ;
```

For reporting purposes, it is useful to define auxiliary variables in `ReferenceModel.py` that will be assigned the cost associated with each stage. These variables do not impact algorithms, but the values are output by some software during execution as well as upon completion. The names of the variables are assigned to stages using the "param StageCostVariable" assignment. The stages are previously defined in `ScenarioStructure.dat` and the variables are previously defined in `ReferenceModel.py`. Note that the variable names appear without the prefix "model."

```
param StageCostVariable := FirstStage  FirstStageCost
                          SecondStage SecondStageCost ;
```

11.2.4 Scenario data specification

So far, we have given a model in the file named `ReferenceModel.py`, a set of deterministic data in the file named `ReferenceModel.dat` and a description of the stochastics in the file named `ScenarioStructure.dat`. All that remains is to give the data for each scenario. There are two ways to do that in PySP: *scenario-based* and *node-based*. The default is scenario-based so we will describe that first.

For scenario-based data, the full data for each scenario is given in a `.dat` file with the root name that is the name of the scenario. So, for example, the file named `AverageScenario.dat` must contain all the data for the model for the scenario named "AverageScenario." It turns out that this file can be created by simply copying the file `ReferenceModel.dat` as shown above because it contains a full set of data for the "AverageScenario" scenario. The files `BelowAverageScenario.dat` and `AboveAverageScenario.dat` will differ from this file and from each other only in their last line, where the yield is specified. These three files are distributed with Coopr and are in the `coopr` sub-directory `examples/pysp/farmer/scenariodata` along with `ScenarioStructure.dat` and `ReferenceModel.dat`.

Scenario-based data wastes resources by specifying the same thing over and over again. In many cases, that does not matter and it is convenient to have full scenario data files available (for one thing, the scenarios can easily be run independently using

the `pyomo` command). However, in many other settings, it is better to use a node-based specification where the data that is unique to each node is specified in a `.dat` file with a root name that matches the node name. In the farmer example, the file `RootNode.dat` will be the same as `ReferenceModel.dat` except that it will lack the last line that specifies the yield. The files `BelowAverageNode.dat`, `AverageNode.dat`, and `AboveAverageNode.dat` will contain only one line each to specify the yield. If node-based data is to be used, then the `ScenarioStructure.dat` file must contain the following line:

```
param ScenarioBasedData := False ;
```

An entire set of files for node-based data for the farmer problem are distributed with Coopr in the sub-directory `examples/pysp/farmer`.

11.3 Finding Solutions for Stochastic Models

PySP provides a variety of tools for finding solutions to stochastic programs.

11.3.1 `runef`

The `runef` command puts together the so-called *extensive form* version of the model. It creates a large model that has constraints to ensure that variables at a node have the same value. For example, in the farmer problem, all of the `DevotedAcres` variables must have the same value regardless of which scenario is ultimately realized. The objective can be the expected value of the objective function, or the CVaR, or a weighted combination of the two. Expected value is the default. A full set of options for `runef` can be obtained using the command:

```
runef --help
```

The `coopr` distribution contains the files need to run the farmer example in the sub-directories to the sub-directory `examples/pysp/farmer` so if this is the current directory and if CPLEX is installed, the following command will cause formation of the EF and its solution using CPLEX.

```
runef -m models -i nodedata --solver=cplex --solve
```

The option `-m models` has one dash and is short-hand for the option `--model-directory=models` and note that the full option uses two dashes. The `-i` is equivalent to `--instance-directory=` in the same fashion. The default solver is CPLEX, so the solver option is not really needed. With the `--solve` option, `runef` would simply write an `.lp` data file that could be passed to a solver.

11.3.2 `runph`

The `runph` command executes an implementation of Progressive Hedging (PH) that is intended to support scripting and extension.

The `coopr` distribution contains the files need to run the farmer example in the sub-directories to the sub-directory `examples/pysp/farmer` so if this is the current directory and if CPLEX is installed, the following command will cause PH to execute using the default sub-problem solver, which is CPLEX.

```
runph -m models -i nodedata
```

The option `-m models` has one dash and is short-hand for the option `--model-directory=models` and note that the full option uses two dashes. The `-i` is equivalent to `--instance-directory=` in the same fashion.

After about 33 iterations, the algorithm will achieve the default level of convergence and terminate. A lot of output is generated and among the output is the following solution information:

```
Variable=DevotedAcreage
  Index: [CORN]          (Scenarios: BelowAverageScenario  AverageScenario  ↔
    AboveAverageScenario )
    Values:      79.9844    80.0000    79.9768    Max-Min=    0.0232  ↔
                Avg=      79.9871
```

```

Index: [SUGAR_BEETS]          (Scenarios: BelowAverageScenario  ↔
AverageScenario  AboveAverageScenario  )
Values:          249.9848    249.9770    250.0000    Max-Min=    0.0230  ↔
                Avg=      249.9873
Index: [WHEAT]              (Scenarios: BelowAverageScenario  AverageScenario  ↔
AverageScenario  AboveAverageScenario  )
Values:          170.0308    170.0230    170.0232    Max-Min=    0.0078  ↔
                Avg=      170.0256
Cost Variable=FirstStageCost
Tree Node=RootNode        (Scenarios: BelowAverageScenario  ↔
AverageScenario  AboveAverageScenario  )
Values:    108897.0836  108897.4725  108898.1476    Max-Min=    1.0640    Avg=  ↔
                108897.5679

```

For problems with no, or few, integer variables, the default level of convergence leaves root-node variables almost converged. Since the acreage to be planted cannot depend on the scenario that will be realized in the future, the average, which is labeled "Avg" in this output, would be used. A farmer would probably interpret acreages of 79.9871, 249.9873, and 170.0256 to be 80, 250, and 170. In real-world applications, PH is embedded in scripts that produce output in a format desired by a decision maker.

But in real-world applications, the default settings for PH seldom work well enough. In addition to post-processing the output, a number of parameters need to be adjusted and sometimes scripting to extend or augment the algorithm is needed to improve convergence rates. A full set of options can be obtained with the command:

```
runph --help
```

Note that there are two dashes before help.

By default, PH uses quadratic objective functions after iteration zero; in some settings it may be desirable to linearize the quadratic terms. This is required to use a solver such as glpk for MIPs because it does not support quadratic MIPs. The directive `--linearize-nonbinary-penalty-terms=n` causes linearization of the penalty terms using `n` pieces. For example, to use glpk on the farmer, assuming glpk is installed and the command is given when the current directory is the `examples/pysp/farmer`, the following command will use default settings for most parameters and four pieces to approximate quadratic terms in sub-problems:

```
runph -i nodedata -m models --solver=glpk --linearize-nonbinary-penalty-terms=4
```

Use of the `linearize-nonbinary-penalty-terms` option requires that all variables not in the final stage have bounds.

11.3.3 Solution Output Control

To get the full solution, including leaf node solution values, use the `runph --output-scenario-tree-solution` option.

In both `runph` and `runef` the solution can be written in csv format using the `--solution-writer=coopr.pysp.csvsolution` option.

11.4 Summary of PySP File Names

PySP scripts such as `runef` and `runph` require files that specify the model and data using files with specific names. All files can be in the current directory, but typically, the file `ReferenceModel.py` is in a directory that is specified using `--model-directory=` option (the short version of this option is `-i +`) and the data files are in a directory specified in the `---instance-directory=` option (the short version of this option is `+m +`).

- `ReferenceModel.py`: A full Pyomo model for a single scenario. There should be no scenario indexes in this model because they are implicit.
- `ReferenceModel.dat`: A full set of data for an arbitrary scenario. This will not be used during solution, but just used to define indexes.

- `ScenarioStructure.dat`: Specifies the nature of the stochastics. It also specifies whether the rest of the data is node-based or scenario-based. It is scenario-based unless `ScenarioStructure.dat` contains the line

```
param ScenarioBasedData := False ;
```

If scenario-based, then there is a data file for each scenario that specifies a full set of data for the scenario. The name of the file is the name of the scenario with `.dat` appended. The names of the scenarios are given in the `ScenarioStructure.dat` file.

If node-based, then there is a file with data for each node that specifies only that data that is unique for the node. The name of the file is the name of the node with `.dat` appended. The names of the nodes are given in the `ScenarioStructure.dat` file.

11.5 Solving Sub-problems in Parallel and/or Remotely

The Python package called Pyro provides capabilities that are used to enable PH to make use of multiple solver processes for sub-problems and allows both `runef` and `runph` to make use remote solvers. We will focus on PH in our discussion here.

There are two solver management systems available for `runph`, one is based on a `pyro_mip_server` and the other is based on a `phsolverserver`. Regardless of which is used, a name server and a dispatch server must be running and accessible to the `runph` process. The name server is launched using the command `coopr_ns` and then the dispatch server is launched with `dispatch_srvr`. Note that both commands contain an underscore. Both programs keep running until terminated by an external signal, so it is common to pipe their output to a file.

Solvers are controlled by solver servers. The `pyro mip solver server` is launched with the command `pyro_mip_server`. This command may be repeated to launch as many solvers as are desired. The `runph` then needs a `--solver-manager=pyro` option to signal that `runph` should not launch its own solver, but should send subproblems to be dispatched to parallel solvers. To summarize the commands:

- Once: `coopr_ns`
- Once: `dispatch_srvr`
- Multiple times: `pyro_mip_server`
- Once: `runph ... --solver-manager=pyro ...`

Note

The `runph` option `--shutdown-pyro` will cause a shutdown signal to be sent to `coopr_ns`, `dispatch_srvr` and all `pyro_mip_server` programs upon termination of `runph`.

Instead of using `pyro_mip_server`, one can use `phsolverserver` in its place. You can get a list of arguments using `pyrosolverserver --help`, which does not launch a solver server (it just displays help and terminates). If you use the `phsolverserver`, then use `--solver-manager=phpyro` as an argument to `runph` rather than `--solver-manager=pyro`.



Warning

Unlike the normal `pyro_mip_server`, there must be one `phsolverserver` for each sub-problem.

Chapter 12

Scripts

There are two main ways to add scripting for Pyomo models: using Python scripts and using callbacks for the `pyomo` command that alter or supplement its workflow.

12.1 Python Scripts

12.1.1 Iterative Example

To illustrate Python scripts for Pyomo we consider an example that is in the file `iterative1.py` and is executed using the command

```
python iterative1.py
```

Note

This is a Python script that contains elements of Pyomo, so it is executed using the `python` command. The `pyomo` command can be used, but then there will be some strange messages at the end when Pyomo finishes the script and attempts to send the results to a solver, which is what the `pyomo` command does.

This script creates a model, solves it, and then adds a constraint to preclude the solution just found. This process is repeated, so the script finds and prints multiple solutions. The particular model it creates is just the sum of four binary variables. One does not need a computer to solve the problem or even to iterate over solutions. This example is provided just to illustrate some elementary aspects of scripting.

Note

The built-in code for printing solutions prints only non-zero variable values. So if you run this code, no variable values will be output for the first solution found because all of the variables are zero. However, other information about the solution, such as the objective value, will be displayed.

```
# iterative1.py
from __future__ import division
from coopr.pyomo import *
from coopr.opt import SolverFactory

# Create a solver
opt = SolverFactory('glpk')
```

```

#
# A simple model with binary variables and
# an empty constraint list.
#
model = AbstractModel()
model.n = Param(default=4)
model.x = Var(RangeSet(model.n), within=Binary)
def o_rule(model):
    return summation(model.x)
model.o = Objective(rule=o_rule)
model.c = ConstraintList()

# Create a model instance and optimize
instance = model.create()
results = opt.solve(instance)
print results

# Iterate to eliminate the previously found solution
for i in range(5):
    instance.load(results)

    expr = 0
    for j in instance.x:
        if instance.x[j].value == 0:
            expr += instance.x[j]
        else:
            expr += (1-instance.x[j])
    instance.c.add( expr >= 1 )

    instance.preprocess()
    results = opt.solve(instance)
    print results

```

Let us now analyze this script. The first line is a comment that happens to give the name of the file. This is followed by two lines that import symbols for Pyomo:

```

# iterativel.py
from coopr.pyomo import *
from coopr.opt import SolverFactory

```

An object to perform optimization is created by calling `SolverFactory` with an argument giving the name of the solver. The argument would be `gurobi` if, e.g., Gurobi was desired instead of `glpk`:

```

# Create a solver
opt = SolverFactory('glpk')

```

The next lines after a comment create a model. For our discussion here, we will refer to this as the base model because it will be extended by adding constraints later. (The words "base model" are not reserved words, they are just being introduced for the discussion of this example). There are no constraints in the base model, but that is just to keep it simple. Constraints could be present in the base model. Even though it is an abstract model, the base model is fully specified by these commands because it requires no external data:

```

model = AbstractModel()
model.n = Param(default=4)
model.x = Var(RangeSet(model.n), within=Binary)
def o_rule(model):
    return summation(model.x)
model.o = Objective(rule=o_rule)

```

The next line is not part of the base model specification. It creates an empty constraint list that the script will use to add constraints.

```
model.c = ConstraintList()
```

The next non-comment line creates the instantiated model and refers to the instance object with a Python variable `instance`. Models run using the `pyomo` script do not typically contain this line because model instantiation is done by the `pyomo` script. In this example, the `create` function is called without arguments because none are needed; however, the name of a file with data commands is given as an argument in many scripts.

```
instance = model.create()
```

The next line invokes the solver and refers to the object contain results with the Python variable `results`.

```
results = opt.solve(instance)
```

The print method of the results object is invoked by the Python `print` command:

```
print results
```

The next non-comment line is a Python iteration command that will successively assign the integers from 0 to 4 to the Python variable `i`, although that variable is not used in script. This loop is what causes the script to generate five more solutions:

```
for i in range(5):
```

The next line associates the results obtained with the instance. This then enables direct queries of solution values in subsequent lines using variable names contained in the instance:

```
    instance.load(results)
```

An expression is built up in the Python variable named `expr`. The Python variable `j` will be iteratively assigned all of the indexes of the variable `x`. For each index, the value of the variable (which was loaded by the `load` method just described) is tested to see if it is zero and the expression in `expr` is augmented accordingly. Although `expr` is initialized to 0 (an integer), its type will change to be a Pyomo expression when it is assigned expressions involving Pyomo variable objects:

```
    expr = 0
    for j in instance.x:
        if instance.x[j].value == 0:
            expr += instance.x[j]
        else:
            expr += (1-instance.x[j])
```

During the first iteration (when `i` is 0), we know that all values of `x` will be 0, so we can anticipate what the expression will look like. We know that `x` is indexed by the integers from 1 to 4 so we know that `j` will take on the values from 1 to 4 and we also know that all value of `x` will be zero for all indexes so we know that the value of `expr` will be something like

```
0 + instance.x[1] + instance.x[2] + instance.x[3] + instance.x[4]
```

The value of `j` will be evaluated because it is a Python variable; however, because it is a Pyomo variable, the value of `instance.x[j]` not be used, instead the variable object will appear in the expression. That is exactly what we want in this case. When we wanted to use the current value in the `if` statement, we used the `value` method to get it.

The next line adds to the constraint list called `c` the requirement that the expression be greater than or equal to one:

```
    instance.c.add( expr >= 1 )
```

The proof that this precludes the last solution is left as an exercise for the reader.

When the model is modified, such as when constraints are added, the `preprocess` method must be called or the changes will not be passed to the solver:

```
    instance.preprocess()
```

The final lines in the outer for loop find a solution and display it:

```
    results = opt.solve(instance)
    print results
```

12.2 Changing the Model or Data and Re-solving

The `iterative.py` example illustrates how a model can be changed and then re-solved. In that example, the model is changed by adding a constraint, but the model could also be changed by altering the values of parameters. Note, however, that in these examples, we make the changes to the `instance` object rather than the `model` object so that we do not have to create a new `model` object. Here is the basic idea:

1. Create an `AbstractModel` (suppose it is called `model`)
2. Call `model.create()` to create an instance (suppose it is called `instance`)
3. Solve `instance`
4. Change something in `instance`
5. Call `presolve`
6. Solve `instance` again

If `instance` has a parameter whose name is in `ParamName` with an index that is in `idx`, the the value in `NewVal` can be assigned to it using

```
getattr(instance, ParamName)[idx] = NewVal
```

For a singleton parameter named `ParamName` (i.e., if it is not indexed), the assignment can be made using either

```
getattr(instance, ParamName)[None] = NewVal
```

or else

```
getattr(instance, ParamName).set_value(NewVal)
```

The function `getattr` is provided by Python. For more information about access to Pyomo parameters, see the section in this document on [Param Access](#).

12.3 Pyomo Callbacks

Pyomo enables altering or extending its workflow through the use of callbacks that are defined in the model file. Taken together, the callbacks allow for construction of a rich set of workflows. However, many users might be interesting in making use of only one or two of the callbacks. They are executable Python functions with pre-defined names:

- `pyomo_preprocess`: Preprocessing before model construction
- `pyomo_create_model`: Constructs and returns the model object
- `pyomo_create_modeldata`: Constructs and returns a `ModelData` object
- `pyomo_print_model`: Display model information
- `pyomo_modify_instance`: Modify the model instance
- `pyomo_print_instance`: Display instance information
- `pyomo_save_instance`: Write the model instance to a file
- `pyomo_print_results`: Display the results of optimization
- `pyomo_save_results`: Store the optimization results
- `pyomo_postprocess`: Postprocessing after optimization

Many of these functions have arguments, which must be declared when the functions are declared. This can be done either by listing the arguments, as we will show below, or by providing a dictionary for arbitrary keyword arguments in the form `**kwds`. If the arbitrary keywords are used, then the arguments are accessed using the `get` method. For example the `preprocess` function takes one argument (as will be described below) so the following two functions will produce the same output:

```
def pyomo_preprocess(options=None):
    if options == None:
        print "No command line options were given."
    else:
        print "Command line arguments were: %s" % options
```

```
def pyomo_preprocess(**kwds):
    options = kwds.get('options',None)
    if options == None:
        print "No command line options were given."
    else:
        print "Command line arguments were: %s" % options
```

To access the various arguments using the `**kwds` argument, use the following strings:

- `options` for the command line arguments dictionary
- `model-options` for the `--model-options` dictionary
- `model` for a model object
- `instance` for an instance object
- `results` for a results object

12.3.1 `pyomo_preprocess`

This function has one argument, which is an enhanced Python dictionary containing the command line options given to launch Pyomo. It is called before model construction so it augments the workflow. It is defined in the model file as follows:

```
def pyomo_preprocess(options=None):
```

12.3.2 `pyomo_create_model`

This function is for experts who want to replace the model creation functionality provided by the `pyomo` script with their own. It takes two arguments: an enhanced Python dictionary containing the command line options given to launch Pyomo and a dictionary with the options given in the `--model-options` argument to the `pyomo` command. The function must return the model object that has been created.

12.3.3 `pyomo_create_modeldata`

Users who employ `ModelData` objects may want to give their own method for populating the object. This function returns a `ModelData` object that will be used to instantiate the model to form an instance. It takes two arguments: an enhanced Python dictionary containing the command line options given to launch Pyomo and a model object.

12.3.4 `pyomo_print_model`

This callback is executed between model creation and instance creation. It takes two arguments: an enhanced Python dictionary containing the command line options given to launch Pyomo and a model object.

12.3.5 `pyomo_modify_instance`

This callback is executed after instance creation. It takes three arguments: an enhanced Python dictionary containing the command line options given to launch Pyomo, a model object, and an instance object.

12.3.6 `pyomo_print_instance`

This callback is executed after instance creation (and after the `pyomo_modify_instance` callback). It takes two arguments: an enhanced Python dictionary containing the command line options given to launch Pyomo and an instance object.

12.3.7 `pyomo_save_instance`

This callback also takes place after instance creation and takes It takes two arguments: an enhanced Python dictionary containing the command line options given to launch Pyomo and an instance object.

12.3.8 `pyomo_print_results`

This callback is executed after optimization. It takes three arguments: an enhanced Python dictionary containing the command line options given to launch Pyomo, an instance object, and a results object. Note that the `--print-results` option provides a way to print results; this callback is intended for users who want to customize the display.

12.3.9 `pyomo_save_results`

This callback is executed after optimization. It takes three arguments: an enhanced Python dictionary containing the command line options given to launch Pyomo, an instance object, and a results object. Note that the `--save-results` option provides a way to store results; this callback is intended for users who want to customize the format or contents.

12.3.10 `pyomo_postprocess`

This callback is also executed after optimization. It also takes three arguments: an enhanced Python dictionary containing the command line options given to launch Pyomo, an instance object, and a results object.

12.4 Accessing Variable Values

12.4.1 Primal Variable Values

Often, the point of optimization is to get optimal values of variables. The `pyomo` script automatically outputs the values to a file and optionally displays the non-zero values on the standard output device (usually the computer screen). Some user may want to process the values in a script. We will describe how to access a particular variable from a Python script as well as how to access all variables from a Python script and from a callback. This should enable the reader to understand how to get the access that they desire. The Iterative example given above also illustrates access to variable values.

12.4.2 One Variable from a Python Script

Assuming the model has been instantiated and solved and the results have been loded back into the instance object, then we can make use of the fact that the variable is a member of the instance object and its value can be accessed using its `value` member. For example, suppose the model contains a variable named `quant` that is a singleton (has no indexes) and suppose further that the name of the instance object is `instance`. Then the value of this variable can be accessed using `instance.quant.value`. Variables with indexes can be referenced by supplying the index.

Consider the following very simple example, which is similar to the iterative example. This is a very simple model and there are no parameter values to be read from a data file, so the `model.create()` call does not specify a file name. In this example, the value of `x[2]` is accessed.

```

# noiteration1.py
from __future__ import division
from coopr.pyomo import *
from coopr.opt import SolverFactory

# Create a solver
opt = SolverFactory('glpk')

#
# A simple model with binary variables and
# an empty constraint list.
#
model = AbstractModel()
model.n = Param(default=4)
model.x = Var(RangeSet(model.n), within=Binary)
def o_rule(model):
    return summation(model.x)
model.o = Objective(rule=o_rule)
model.c = ConstraintList()

# Create a model instance and optimize
instance = model.create()
results = opt.solve(instance)
instance.load(results)

if instance.x[2].value == 0:
    print "The second index has a zero"
else:
    print "x[2]=", instance.x[2].value

```

12.4.3 All Variables from a Python Script

As with one variable, we assume that the model has been instantiated and solved and the results have been loaded back into the instance object using `instance.load(results)`, then we can make use of the fact that the variable is a member of the instance object and its value can be accessed using its value member. Assuming the instance object has the name `instance`, the following code snippet displays all variables and their values:

```

from coopr.pyomo import Var
for v in instance.active_components(Var):
    print "Variable", v
    varobject = getattr(instance, v)
    for index in varobject:
        print "    ", index, varobject[index].value

```

This code could be improved by checking to see if the variable is not indexed (i.e., the only index value is `None`), then the code could print the value without the word `None` next to it.

12.4.4 All Variables from Workflow Callbacks

The `pyomo_print_results`, `pyomo_save_results`, and `pyomo_postprocess` callbacks from the `pyomo` script take the instance as one of their arguments and the instance has the solver results at the time of the callback so the body of the callback matches the code snippet given for a Python script.

For example, if the following definition were included in the model file, then the `pyomo` command would output all variables and their values (including those variables with a value of zero):

```

def pyomo_print_results(options, instance, results):
    from coopr.pyomo import Var

```

```

for v in instance.active_components(Var):
    print "Variable",v
    varobject = getattr(instance, v)
    for index in varobject:
        print "    ",index, varobject[index].value

```

12.5 Accessing Parameter Values

Access to parameters is completely analogous to access to variables. For example, here is a code snippet to print the name and value of every Parameter:

```

from coopr.pyomo import Param
for p in instance.active_components(Param):
    print "Parameter",p
    parmobject = getattr(instance, p)
    for index in parmobject:
        print "    ",index, parmobject[index].value

```

12.6 Accessing Duals

Access to dual values in scripts is similar to accessing primal variable values, except that dual values are not captured by default so additional directives are needed before optimization to signal that duals are desired.

To get duals without a script, use the `pyomo` option `--solver-suffixes=dual` which will cause dual values to be included in output. Note: In addition to duals (`dual`), reduced costs (`rc`) and slack values (`slack`) can be requested. All suffixes can be requested using the `pyomo` option `--solver-suffixes=.*`

Warning: Some of the duals may have the value `None`, rather than `0`.

12.6.1 Access Duals in a Python Script

To signal that duals are desired, add the argument `suffixes=['dual']` to the `opt.solve` function call. After the results are obtained, duals can be accessed in a fashion analogous to access of primal variable values.

```

# display all duals
print "Duals"
from coopr.pyomo import Constraint
for c in instance.active_components(Constraint):
    print "  Constraint",c
    cobject = getattr(instance, c)
    for index in cobject:
        print "    ",index, cobject[index].dual

```

The following snippet will only work, of course, if there is a constraint with the name `AxbConstraint` that has an index, which is the string `Film`.

```

# access (display, this case) one dual
print "Dual for Film=", instance.AxbConstraint['Film'].dual

```

Here is a complete example that relies on the file `abstract2.py` to provide the model and the file `abstract2.dat` to provide the data. Note that the model in `abstract2.py` does contain a constraint named `AxbConstraint` and `abstract2.dat` does specify an index for it named `Film`.

```

# driveabs2.py
from coopr.pyomo import *
from __future__ import division
from coopr.opt import SolverFactory

# Create a solver
opt = SolverFactory('cplex')

# get the model from another file
from abstract2 import model

# Create a model instance and optimize
instance = model.create('abstract2.dat')
results = opt.solve(instance, suffixes=['dual'])

# get the results back into the instance for easy access
instance.load(results)

# display all duals
print "Duals"
from coopr.pyomo import Constraint
for c in instance.active_components(Constraint):
    print "  Constraint",c
    cobject = getattr(instance, c)
    for index in cobject:
        print "      ", index, cobject[index].dual

# access (display, this case) one dual
print "Dual for Film=", instance.AxbConstraint['Film'].dual

```

12.6.2 All Duals from Workflow Callbacks

The `pyomo` script needs to be instructed to obtain duals, either by using a command line option such as `--solver-suffixes=.dual` or by adding code in the `pyomo_preprocess` callback to add `solver-suffixes` to the list of command line arguments if it is not there and to add `.dual` to its list of arguments if it is there, but `.dual` is not.

The `pyomo_print_results`, `pyomo_save_results`, and `pyomo_postprocess` callbacks from the `pyomo` script take the instance as one of their arguments and the instance has the solver results at the time of the callback so the body of the callback matches the code snipped given for a Python script.

For example, if the following definition were included in the model file, then the `pyomo` command would output all constraints and their duals.

```

def pyomo_print_results(options, instance, results):
    # display all duals
    print "Duals"
    from coopr.pyomo import Constraint
    for c in instance.active_components(Constraint):
        print "  Constraint",c
        cobject = getattr(instance, c)
        for index in cobject:
            print "      ", index, cobject[index].dual

```

12.7 Accessing Solver Status

After a solve, the results object has a member `Solution.Status` that contains the solver status. The following snippet shows an example of access via a `print` statement:

```
instance = model.create()
results = opt.solve(instance)
print "The solver returned a status of:"+str(results.Solution.Status)
```

The use of the Python `str` function to cast the value to a be string makes it easy to test it. In particular, the value *optimal* indicates that the solver succeeded.

12.8 Display of Solver Output

To see the output of the solver, use the option `tee=True` as in

```
results = opt.solve(instance, tee=True)
```

This can be useful for troubleshooting solver difficulties.

Chapter 13

Coopr Solver Interfaces

This chapter describes how Coopr interfaces with different solvers.

Chapter 14

Using Black-Box Optimizers with Coopr.Opt

Many optimization software packages contain *black-box* optimizers, which perform optimization without using detailed knowledge of the structure of an optimization problem. Thus, black-box optimizers require a generic interface for optimization problems that defines key features of problems, like objectives and constraints.

The `coopr.opt` package contains the `coopr.opt.blackbox` subpackage, which provides facilities for (a) integrating Coopr solvers with blackbox optimization applications and (b) wrapping Pyomo models for use by external blackbox optimizers. We illustrate these capabilities in this chapter with simple examples that illustrate the use of `coopr.opt.blackbox`.

14.1 Defining and Optimizing Simple Black-Box Applications

Many black-box optimizers interact with an optimization problem by executing a separate process that computes properties of the optimization problem. This process typically reads an input file that defines the requested properties and writes an output file that contains the computed values. Unfortunately, no standards have emerged for black-box optimizers that interact with problems in this manner. Thus, different file formats are used by different optimizer software packages.

14.1.1 Defining an Optimization Problem

The `coopr.opt.blackbox` package provides several Python classes for optimization problems that coordinates file I/O for the user and simplifies the definition of simple black-box problems. The `RealOptProblem` class provides a generic interface for continuous optimization problems (i.e. with real variables). The following example defines a simple continuous optimization problem:

```
class RealProblem1(RealOptProblem):

    def __init__(self):
        RealOptProblem.__init__(self)
        self.lower=[0.0, -1.0, 1.0, None]
        self.upper=[None, 0.0, 2.0, -1.0]
        self.nvars=4

    def function_value(self, point):
        self.validate(point)
        return point.vars[0] - point.vars[1] + (point.vars[2]-1.5)**2 + (point.vars[3]+2) ←
            **4
```

This problem is equivalent to the following problem definition:

$$\begin{aligned} \min \quad & x_0 - x_1 + (x_2 - 1.5)^2 + (x_3 + 2)^4 \\ \text{s.t.} \quad & 0 \leq x_0 \\ & -1 \leq x_1 \leq 0 \\ & 0 \leq x_2 \leq 2 \\ & x_3 \leq -1 \end{aligned}$$

Note that the problem class does *not* specify the sense of the optimization problem. These problem classes are not a complete specification of an optimization problem. Rather, an instance of a problem class can compute information about the problem that is used during optimization.

Similarly, the `MixedIntOptProblem` class provides a generic interface for mixed-integer optimization problems, which may contain real variables, integer variables and binary variables. The following example defines a simple mixed-integer optimization problem:

```
class MixedIntProblem1(MixedIntOptProblem):

    def __init__(self):
        MixedIntOptProblem.__init__(self)
        self.real_lower=[0.0]*4
        self.real_upper=[2.0]*4
        self.int_lower=[-2]*3
        self.int_upper=[0]*3
        self.nreal=4
        self.nint=3
        self.nbinary=2

    def function_value(self, point):
        self.validate(point)
        return sum((x-1)**2 for x in self.reals) + \
            sum((y+1)**2 for y in self.ints) + \
            sum(b for b in self.bits)
```

This problem is equivalent to the following problem definition:

$$\begin{aligned} \min \quad & \sum_{i=1}^4 (x_i - 1)^2 + \sum_{i=1}^3 (y_i + 1)^2 + \sum_{i=1}^2 z_i \\ \text{s.t.} \quad & 0 \leq x_i \leq 2 \\ & -2 \leq y_i \leq 0 \\ & z_i \in \{0, 1\} \end{aligned}$$

14.1.2 Optimizing with Coliny Solvers

The `Coliny` software library supports interfaces to a variety of black-box optimizers <Coliny>. The `coliny` executable reads an XML specification of the optimization problem and solver, as well as a specification of how the optimizer is applied. Consider the following XML specification:

```
<!-- RealProblem1.xml

    This Coliny XML specification illustrates the execution of the
    colin:ls solver on the RealProblem1 problem.
-->

<ColinInput>

  <Problem type="MINLP0">
    <Domain>
      <RealVars num="4">
        <Lower index="1" value="0.0"/>
        <Lower index="2" value="-1.0"/>
        <Lower index="3" value="1.0"/>

        <Upper index="2" value="0.0"/>
        <Upper index="3" value="2.0"/>
        <Upper index="4" value="-1.0"/>
      </RealVars>
    </Domain>

    <Driver>
```

```

    <Command>RealProblem1.py</Command>
  </Driver>
</Problem>

<Solver type="colin:ls">
  <InitialPoint>
    0.0 2.0 -1.0 10.0
  </InitialPoint>
  <Options>
    <Option name="min_function_value">-1.0</Option>
  </Options>
</Solver>
</ColinInput>

```

This XML specification defines a MINLP0 problem, which indicated that this is a mixed-integer problem that supports zero-order derivatives (i.e. no derivatives). This problem has four real variables with lower and upper bounds specified. The problem values are computed with the `RealProblem1.py` command-line, which defines and uses the `RealProblem1` class defined above:

```

#!/usr/bin/env python
#
# RealProblem1.py

import sys
from coopr.opt.blackbox import RealOptProblem

class RealProblem1(RealOptProblem):

    def __init__(self):
        RealOptProblem.__init__(self)
        self.lower=[0.0, -1.0, 1.0, None]
        self.upper=[None, 0.0, 2.0, -1.0]
        self.nvars=4

    def function_value(self, point):
        self.validate(point)
        return point.vars[0] - point.vars[1] + (point.vars[2]-1.5)**2 + (point.vars[3]+2) ←
            **4

problem = RealProblem1()
problem.main(sys.argv)

```

Note that this command is a Python script that includes the shebang character sequence on the first line. On Linux and Unix systems, this line indicates that this is a script that is executed using the `python` command that is found in the user environment. Thus, this example assumes that the `python` command has `coopr.opt` installed. Since multiple versions of Python can be installed on a single computer, the XML Command element may need to be defined with an explicitly Python version. For example, if Python 2.6 is installed in `/usr/local` with `coopr.opt`, then the Command element would look like:

```

<Command>/usr/local/bin/python26 RealProblem1.py</Command>

```

Additionally, the duplication of bounds information between `RealProblem1.py` and `RealProblem1.xml` is not strictly necessary in this example. The bounds information in `RealProblem1.py` is used in the `validate` method to verify that the point being evaluated is consistent with the bounds information. We can generally assume that the Coliny solver will only evaluate feasible points, so a simpler problem definition can be used:

```

#!/usr/bin/env python
#
# RealProblem2.py

import sys
from coopr.opt.blackbox import RealOptProblem

```

```

class RealProblem2(RealOptProblem):

    def __init__(self):
        RealOptProblem.__init__(self)
        self.nvars=4

    def function_value(self, point):
        return point.vars[0] - point.vars[1] + (point.vars[2]-1.5)**2 + (point.vars[3]+2) ←
            **4

problem = RealProblem2()
problem.main(sys.argv)

```

The last two lines of `RealProblem1.py` create a problem instance and then call the `main` method to parse the command-line arguments. This script has the following command-line syntax:

```
RealProblem1.py <input-file> <output-file>
```

The first argument is the name of an XML input file, and the second argument is the name of an XML output file. The optimization problem class manages the parsing of the input and generation of the output file. For example, consider the following input file:

```

<ColinRequest>
  <Parameters>
    <Real size="4"> 0.1e-1 -0.1 1.1 -1.9</Real>
  </Parameters>
  <Requests>
    <FunctionValue/>
  </Requests>
</ColinRequest>

```

The `RealProblem1.py` script creates the following output file:

```

<?xml version="1.0" encoding="UTF-8"?>
<ColinResponse>
  <FunctionValue>
    0.2701
  </FunctionValue>
</ColinResponse>

```

14.2 Diving Deeper

The previous section provided an overview of the how the `coopr.opt.blackbox` package supports the definition of optimization problems that are solved with black-box optimizers. In this section we provide more detail about how the Python problem class can be customized, as well as details about the XML file format used to communicate with Coliny optimizers. The Dakota User Manual <Dakota> provides documentation of the file format of the input and output files used with Dakota optimizers.

Table Table 14.1 summarizes the methods of the `OptProblem` class that a user is likely to either use or redefine when declaring a subclass. The `MixedIntOptProblem` class is a convenient base class for the problems solved by most black-box optimizers, and this class provides the definition of the `main`, `create_point` and `validate` methods. However, any of the remaining methods may need to be defined, depending on the problem.

Table 14.1: Methods in the `OptProblem` class.

Method	Description
<code>__init__</code>	The constructor, which may be redefined to specify problem properties.

Table 14.1: (continued)

main	Method that processes command-line options to create a results file from an input file.
create_point	Create an instances of the class that defines a point in the search domain.
function_value	Returns the value of the objective function.
function_values	Returns a list of objective function values.
gradient	Returns a list that represents the gradient vector at the given point.
hessian	Returns a Hessian matrix.
nonlinear_constraint_values	Returns a list of values for the constraint functions.
jacobian	Returns a Jacobian matrix.
validate	Returns True if the given point is feasible, and False otherwise.

The following detailed example illustrates the use of all of these methods in a simple application:

```
class RealProblem3(RealOptProblem):

    def __init__(self):
        RealOptProblem.__init__(self)
        self.nvars=4
        self.ncons=4
        self.response_types = [response_enum.FunctionValue,
                               response_enum.Gradient,
                               response_enum.Hessian,
                               response_enum.NonlinearConstraintValues,
                               response_enum.Jacobian]

    def function_value(self, point):
        return point.vars[0] - point.vars[1] + (point.vars[2]-1.5)**2 + (point.vars[3]+2) ←
            **4

    def gradient(self, point):
        return [1, -1, 2*(point.vars[2]-1.5), 4*(point.vars[3]+2)**3]

    def hessian(self, point):
        H = []
        H.append( (2,2,2) )
        H.append( (3,3,12*(point.vars[3]+2)**2) )
        return H

    def nonlinear_constraint_values(self, point):
        C = []
        C.append( sum(point.vars) )
        C.append( sum(x**2 for x in point.vars) )
        return C

    def jacobian(self, point):
        J = []
        for j in range(self.nvars):
            J.append( (0,j,1) )
        for j in range(self.nvars):
            J.append( (1,j,2*point.vars[j]) )
        return J
```

The `response_types` attribute defined in the constructor specifies the type of information that this class can compute. For example, consider the following input XML file:

```

<ColinRequest>
  <Parameters>
    <Real size="4"> 0.1e-1 -0.1 1.1 -1.9</Real>
  </Parameters>
  <Requests>
    <FunctionValue/>
    <Gradient/>
    <Hessian/>
    <NonlinearConstraintValues/>
    <Jacobian/>
  </Requests>
</ColinRequest>

```

This input file requests that the class compute all of the response values, and thus the following output is generated:

```

<?xml version="1.0" encoding="UTF-8"?>
<ColinResponse>
  <Gradient>
    1 -1 -0.79999999999999982 0.00400000000000000105
  </Gradient>
  <NonlinearConstraintValues>
    -0.88999999999999999 4.8300999999999998
  </NonlinearConstraintValues>
  <FunctionValue>
    0.2701
  </FunctionValue>
  <Hessian>
    (2, 2, 2) (3, 3, 0.120000000000000022)
  </Hessian>
  <Jacobian>
    (0, 0, 1) (0, 1, 1) (0, 2, 1) (0, 3, 1) (1, 0, 0.02) (1, 1, -0.20000000000000001) (1, 2, ←
    2.2000000000000002) (1, 3, -3.7999999999999998)
  </Jacobian>
</ColinResponse>

```

Note that the values for Jacobian and Hessian matrices are represented in a sparse manner. Currently, these are represented with a list of tuple values, though a sparse matrix representation might be supported in the future.

Chapter 15

Distributed Optimization with Pyro

Coopr supports distributed computing via the Python "PYRO" package. PYRO stands for PYthon Remote Objects. Full documentation of PYRO is available from: <http://pyro.sourceforge.net/>.

The following describes a "quick-start" process for creating and using a client and multiple solvers on a single, presumably multi-core compute server. For example, an institution may have an 8-core workstation with numerous CPLEX licenses. With distributed solves under PYRO, Coopr algorithms can take advantage of the full set of resources on a machine.

The following example assumes a unix/linux platform. The steps for Windows are qualitatively identical - the sole difference is that you can't (or at least we haven't figured out how to) put processes in the background on Windows. The work-around is simply (albeit painfully) to launch the various processes in distinct shells.

15.1 Step 1: Starting a Name Server

All PYRO objects communicate via a name server, which provides a well-defined point of contact through which distributed objects can interact. You can think of the name server as a phone directory.

To start the name server, type:

```
coopr-ns
```

In general, we suggest that the output be redirected to a file, with the entire process being placed in the background:

```
coopr-ns >& ns.out &
```

15.2 Step 2: Starting a Dispatch Server

With the name server up and running, the next step is to create a dispatch server. The function of the dispatch server is to route work from clients to servers - both of the latter will be established in the immediately following steps. We again assume the process is executed in the background, with the output redirected:

```
dispatch_srvr >& dispatch_srvr.out &
```

15.3 Step 3: Starting a MIP server

With the work dispatcher up, the next step is to create servers to do real work! Coopr ships with a `pyro_mip_server` script, which launches a server capable of solving a single MIP at a time. This server can be invoked as follows:

```
pyro_mip_server >& pyro_mip_server1.out &
```

We can also create multiple instances of the `pyro_mip_server`, e.g., to take advantage of multiple solver licenses:

```
pyro_mip_server >& pyro_mip_server2.out &
```

With this configuration, the dispatch server "sees" two mip servers, and can route work to both.

15.4 Step 4: Running a Client

To take advantage of the distributed MIP servers, a Coopr user only needs to change the type of the solver manager supplied to the various solver scripts.

For example, one can run `pyomo` as follows, considering the PySP example found in: `coopr/examples/pysp/farmer`:

```
pyomo --solver=cplex --solver-manager=pyro farmer_lp.py farmer_lp.dat
```

This will execute the LP solve using one of the two mip servers established in Step 3.

To take advantage of parallelism, we can solve the farmer example using progressive hedging, as follows:

```
runph --solver=cplex --solver-manager=pyro --model-directory=models --instance-directory= ↵  
scenariodata
```

15.5 Moving from Multi-Core to Distributed Computation

Truly distributed computation, i.e., with the client and server components on different hosts, is only incrementally more difficult than what is outlined above. If multiple hosts are involved in the computation, the only real issue is making sure the various hosts can all find a common nameserver. After starting `coopr-ns` on some host (presumably a server-class machine), the other components (`dispatch_srvr` and the `pyro_mip_server`) can be pointed to the nameserver by simply setting the environment variable `PYRO_NS_HOSTNAME` to the name (or IP address) of the host running the nameserver. The same process should be followed on the client prior to executing either `pyomo`, `runph`, or some other client solver script.

We have tested this on linux clusters with success. The only issues encountered involve overly aggressive firewalls on the host running the nameserver, which was easily corrected. In theory, Pyro should also work on Windows clusters, and linux-Windows hybrid clusters via the same mechanism.

15.6 Cleaning Up After Yourself

It is important to remember that the name server, the dispatch server, and the mip server processes are persistent, and need to be terminated when a user has completed computational experiments. Actually, that is not entirely correct - the server processes can live forever, and continue to receive work. The only issue is when multiple users are attempting to use the same compute platform, are running their own servers, etc. While this may work, we have not tested it fully yet!

Bibliography

- [AIMMS] <http://www.aimms.com/>
 - [AMPL] <http://www.ampl.com>
 - [GAMS] <http://www.gams.com>
 - [Coliny] Hart and Siirola. Coliny Technical Report. Sandia National Labs.
 - [Dakota] Eldred et al. Dakota Technical Report. Sandia National Labs.
 - [MPG] Mathematical Programming Glossary, <http://glossary.computing.society.informs.org/>, Ed. Allen Holder, 2012.
 - [PyomoBook] Hart, W.E., Laird, C., Watson, J.-P., Woodruff, D.L., Pyomo – Optimization Modeling in Python, Springer, 2012.
 - [PyomoJournal] William E. Hart, Jean-Paul Watson, David L. Woodruff, "Pyomo: modeling and solving mathematical programs in Python," Mathematical Programming Computation, Volume 3, Number 3, August 2011
 - [BirgeLouveauxBook] Introduction to Stochastic Programming; J.R. Birge and F. Louveaux; Springer Series in Operations Research; New York: Springer, 1997.
 - [Vielma_et_al] Vielma, J.P, Ahmed, S., Nemhauser, G., "Mixed-Integer Models for Non-separable Piecewise Linear Optimization: Unifying framework and Extensions," Operations Research 58, 2010. pp. 303-315.
-

Colophon

This book was created using `asciidoc` software.

The Sandia National Laboratories Laboratory-Directed Research and Development Program and the U.S. Department of Energy Office of Science Advanced Scientific Computing Research Program funded portions of this work.

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND 2012-1572P.

Sandia National Laboratories

United States Department of Energy
