
AIMMS User's Guide - Case Management

This file contains only one chapter of the book. For a free download of the complete book in pdf format, please visit www.aimms.com or order your hard-copy at www.lulu.com/aimms.

Copyright © 1993–2011 by Paragon Decision Technology B.V. All rights reserved.

Paragon Decision Technology B.V.	Paragon Decision Technology Inc.	Paragon Decision Technology Pte.
Schipholweg 1	500 108th Avenue NE	Ltd.
2034 LS Haarlem	Ste. # 1085	80 Raffles Place
The Netherlands	Bellevue, WA 98004	UOB Plaza 1, Level 36-01
Tel.: +31 23 5511512	USA	Singapore 048624
Fax: +31 23 5511517	Tel.: +1 425 458 4024	Tel.: +65 9640 4182
	Fax: +1 425 458 4025	

Email: info@aimms.com
WWW: www.aimms.com

AIMMS is a registered trademark of Paragon Decision Technology B.V. IBM ILOG CPLEX and sc CPLEX is a registered trademark of IBM Corporation. GUROBI is a registered trademark of Gurobi Optimization, Inc. KNITRO is a registered trademark of Ziena Optimization, Inc. XPRESS-MP is a registered trademark of FICO Fair Isaac Corporation. MOSEK is a registered trademark of Mosek ApS. WINDOWS and EXCEL are registered trademarks of Microsoft Corporation. $\text{T}_{\text{E}}\text{X}$, $\text{L}_{\text{A}}\text{T}_{\text{E}}\text{X}$, and $\text{A}_{\text{M}}\text{S}_{\text{L}}\text{A}_{\text{T}}\text{E}_{\text{X}}$ are trademarks of the American Mathematical Society. LUCIDA is a registered trademark of Bigelow & Holmes Inc. ACROBAT is a registered trademark of Adobe Systems Inc. Other brands and their products are trademarks of their respective holders.

Information in this document is subject to change without notice and does not represent a commitment on the part of Paragon Decision Technology B.V. The software described in this document is furnished under a license agreement and may only be used and copied in accordance with the terms of the agreement. The documentation may not, in whole or in part, be copied, photocopied, reproduced, translated, or reduced to any electronic medium or machine-readable form without prior consent, in writing, from Paragon Decision Technology B.V.

Paragon Decision Technology B.V. makes no representation or warranty with respect to the adequacy of this documentation or the programs which it describes for any particular purpose or with respect to its adequacy to produce any particular result. In no event shall Paragon Decision Technology B.V., its employees, its contractors or the authors of this documentation be liable for special, direct, indirect or consequential damages, losses, costs, charges, claims, demands, or claims for lost profits, fees or expenses of any nature or kind.

In addition to the foregoing, users should recognize that all complex software systems and their documentation contain errors and omissions. The authors, Paragon Decision Technology B.V. and its employees, and its contractors shall not be responsible under any circumstances for providing information or corrections to errors and omissions discovered at any time in this book or the software it describes, whether or not they are aware of the errors or omissions. The authors, Paragon Decision Technology B.V. and its employees, and its contractors do not recommend the use of the software described in this book for applications in which errors or omissions could threaten life, injury or significant loss.

This documentation was typeset by Paragon Decision Technology B.V. using $\text{L}_{\text{A}}\text{T}_{\text{E}}\text{X}$ and the LUCIDA font family.

Part IV

Data Management

Chapter 17

Case Management

Case management is an important part of any successful decision support application. The ability to save and work with many cases, and perform a what-if analysis by computing and analyzing the results for a large number of scenarios helps in taking the right decisions in real-life situations. This chapter introduces AIMMS' capabilities with respect to creating and managing a large database of cases, as well as its capabilities of working with data from multiple cases both within the language and in graphical data objects on end-user pages.

This chapter

17.1 What are cases?

A case forms a *complete* set of data that will enable you to restart a particular modeling application. In the simplest form of case management, which will be the subject of this chapter, a case contains the data associated with *all* identifiers in your model.

What are cases?

For more advanced use, AIMMS also supports the concepts of

Advanced use

- *case types*, with which you can instruct AIMMS to only save the data of a subset of the identifiers in your model, and
- *data categories* and *datasets*, which allow you to store common data, that is shared by multiple cases, at a single location.

Both of these subjects are discussed in [Chapter 18](#).

In this chapter, you will find only a single case type, *All Identifiers*, which results in cases containing the data of all the identifiers in your model. This case type is automatically added to every new AIMMS project, and all dialog boxes referring to case types will default to it.

All Identifiers

When you want to work with cases in AIMMS, there are two main tools through which you can accomplish most tasks. These are:

Case management tasks

- the **Data** menu, through which you can accomplish simple case management tasks such as loading and saving cases from within your modeling application, and

- the **Data Manager**, which you can use to manage the complete collection of cases in your application, to create batch runs of cases, or to create a selection of cases for simultaneous display.

The following two sections will discuss both tools in more detail.

17.2 Working with cases

While the **Data Manager** is mainly intended for organizing and managing a large collection of datasets and cases, common end-user case management tasks such as loading case data into your model, and saving the current data in your model into cases are usually performed using the **Data** menu only. By default, the **Data** menu is available on all end-user pages.

Loading and saving

In AIMMS, all the data that you are currently working with is referred to as the *active case*. If you have not yet loaded or saved a case, the active case is *unnamed*, otherwise the active case is *named* after the name of the last loaded or saved case on disk. If the active case is named, its name is displayed in the status bar at the bottom of the AIMMS window.

The active case

When you save a named active case, AIMMS will save it to the associated case on disk by default (thus overwriting its previous contents). If the active case is unnamed, or when you try to save a case using the **Data-Save Case As** menu, AIMMS will open the **Save Case** dialog box illustrated in Figure 17.1. In the

Saving a case




Figure 17.1: The **Save Case** dialog box

Save Case dialog box you can enter the name of the case, and, optionally, select the folder in the case tree (explained below) in which the case is to be

stored. After successfully saving a case through the **Save Case** dialog box, the active case will become named.

AIMMS supports three modes for loading the data of a case:

Loading a case

- load as active,
- load into active, and
- merge into active.

These three modes of loading a case differ in

- whether they change the name of the active case or are only importing data into your current active case, and
- whether existing data is replaced by, or merged with, the loaded data.

The most frequently used mode for loading a case is loading the case *as active*, through the **Data-Load Case-As Active** menu. Loading a case as active completely replaces the active case data of all identifiers in the loaded case with their stored values. Data of identifiers that are not stored in the case, remain unchanged. In addition, the active case will be named after the loaded case. Before loading a case as active, AIMMS will ask you whether the current active case data needs to be saved whenever this is necessary.

Load as active

Loading a case *into active* (through the **Data-Load Case-Into Active** menu) is completely identical to loading a case as active, with the exception that the name of the active case will not be changed. Thus, by loading data into the active case you can replace part, or all, of the contents of the active case with data obtained from another case.

Load into active

Merging a case into active (through the **Data-Load Case-Merge Into Active** menu) does not change the name of the active case either. Merging a case into active partially replaces the data in the active case with only the nondefault values stored in the loaded case. Data in the active case, for which no associated nondefault values exist in the loaded case, remain unchanged.

Merge into active

Using the **Data-New Case** menu item, you can instruct AIMMS to start a new, unnamed, active case. However, the data in the active case will remain *unchanged*. If you also want to remove all data from the active case, you can accomplish this from within your model using the `EMPTY` statement. Before starting a new case, AIMMS will ask you whether the current active case data needs to be saved.

Starting a new case

17.3 Managing cases with the Data Manager

The management of all cases (and also datasets) stored within a particular AIMMS project is done through the AIMMS **Data Manager**. It offers you a tree-based view of all the cases created in your application, as illustrated in Figure 17.2. The figure displays a subset of the entire collection of cases. The node name of the current active case is displayed in bold.

The Data Manager

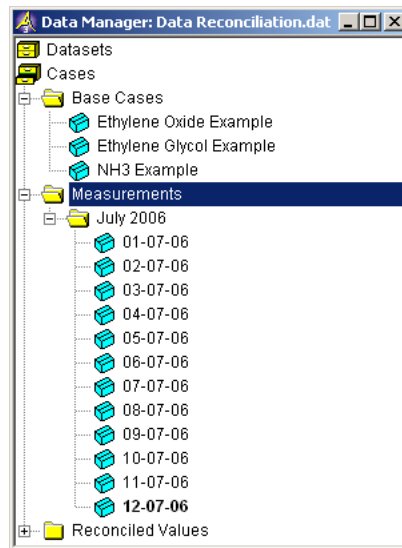


Figure 17.2: The AIMMS Data Manager

Below the case tree in the **Data Manager**, you can add any hierarchical sub-structure of folders to organize the available cases as you see fit. This allows you, for example, to subdivide the entire collection of cases by user or by any other criteria that you, or your end-users, deem relevant. Within this hierarchical structure of folders you can insert, copy and delete cases, or move cases around in the same manner as in any of the other AIMMS trees (see also Section 4.3).

Managing cases

In addition to creating cases through the **Data** menu, you can also create new cases in the **Data Manager** itself. Creating cases in this manner can simply be accomplished by adding a node to the case tree, in a similar fashion as with any other tree. When you create a case in the **Data Manager**, it will not contain any data initially.

Creating cases in the Data Manager

Creating cases in the **Data Manager** is particularly useful when you want to quickly compose a large number of scenarios which all make use of a number of datasets containing common data shared by all scenarios. Including datasets in a case created in the **Data Manager** is discussed in full detail in Section 18.3.

Including datasets

Within the **Data Manager**, you can view or modify the properties of all existing cases through the case **Properties** dialog box illustrated in Figure 17.3. This

Viewing the case properties




Figure 17.3: The case **Properties** dialog box

dialog box lets you provide further identification information for the case, such as an acronym for use in a multiple-case display (see Section 17.5.1) and an explanatory comment. In addition, in the **Data** tab, you can view or modify the case type (see Section 18.1) associated with the case, as well as view the datasets included in the case (see also Section 18.2.1).

17.4 Executing batch runs

The case facilities in AIMMS make it very easy to perform a what-if analysis on your model. When you store the input data for each scenario in a separate case, performing a what-if analysis boils down to running your model for a batch of cases, and comparing the results stored in these cases.

What-if analysis

After you have created a number of cases containing the (input) scenarios for a what-if analysis, you need to compute the optimal solution for all of these scenarios. If you have only a few scenarios which solve relatively quickly, you can perform this interactively through the following sequence of actions:

Interactive batch

- load the case data,

- execute the procedure that computes the model results, and
- save the results back to the case.

If there are many scenarios to be solved, or if the solution time of each individual scenario is long, AIMMS offers facilities to perform the actual what-if analysis by creating and executing a *batch run* of scenarios. Such a batch run can be executed over night, or at any other time that you do not need your computer.

Batch runs

You can add one or more selected cases from within the **Data Manager** to a batch run through the **Edit-Add To-Batch** menu in the **Data Manager**. This will open the **Batch Run** window illustrated in Figure 17.4. This window shows the

Adding cases to a batch run

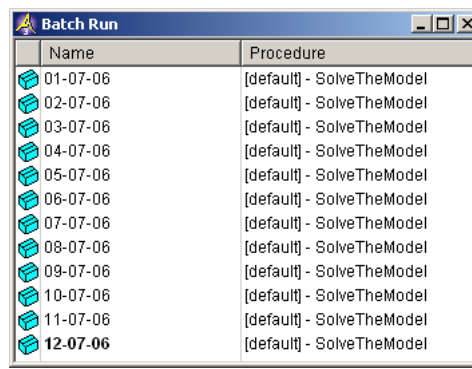


Figure 17.4: The **Batch Run** window

cases that have been already added to the batch run, along with the procedures within the model that must be executed for each case in the batch run to compute the model results.

Within the **Batch Run** window, you can further manipulate the current list of cases. You can modify the order of execution by changing the case order in the list of cases, delete cases from the list, or add cases to the batch run by dragging them from within the **Data Manager** into the **Batch Run** window.

Managing a batch run

In addition, the **Batch Run** window lets you specify a specific procedure for each case in the list that has to be executed to compute the model results for that particular case in the batch run. You can modify the batch procedure for a case through the **Run-Procedure** menu, which will let you select a procedure from the list of all available procedures in the model.

Assigning a procedure

If you do not specify a specific procedure to be executed for a particular case, AIMMS will execute a default procedure. As the developer of a project, you can specify this default procedure through the **Run-Default Procedure** menu, and AIMMS stores this information along with the project. End-users are not allowed to modify the default batch procedure. In the **Batch Run** window, AIMMS displays the default batch procedure for every case for which you have not selected a case-specific batch procedure.

The default batch procedure

After you have composed a batch of cases to your satisfaction, you can start the execution of the batch through the **Run-Start Batch** menu. This will open the dialog box shown in Figure 17.5. It displays the total number of cases in

Starting a batch

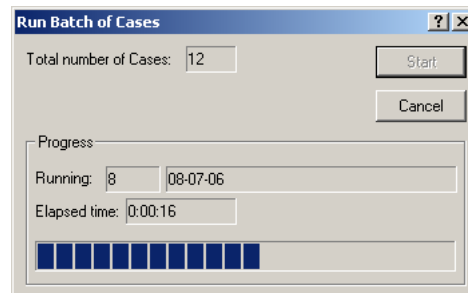


Figure 17.5: The **Run Batch** dialog box

the batch, some information about the case currently being executed, as well as the overall progress of the batch run. By pressing the **Cancel** button in the dialog box, you can interrupt the execution of a running batch.

During a batch run, AIMMS will perform the following actions for every case in the batch run:

Executing a single case

- load the case,
- run the specified or default batch procedure, and
- save the case in order to store the model results.

17.5 Managing multiple case selections

After you have executed a batch run (or when you have created several cases manually), AIMMS allows you to simultaneously view the results of several cases within the graphical user interface. In addition, it is possible to reference data from multiple cases within the modeling language, enabling you to perform advanced forms of case comparison.

Viewing batch results

AIMMS offers a special window, similar to the **Batch Run** window, to construct a selection of cases to which you want simultaneous access either from within the graphical user interface or from within the model itself. You can add one or more selected cases from within the **Data Manager** to the multiple case selection through the **Edit-Add To-Multiple Cases** menu in the **Data Manager**. This will open the **Multiple Cases** window illustrated in Figure 17.6. It shows

Multiple case selections

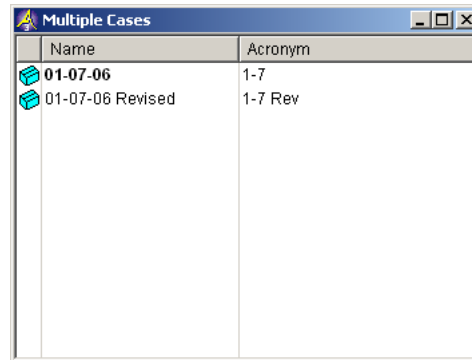


Figure 17.6: The **Multiple Cases** window

the current contents of the multiple case selection. As with the **Batch Run** window, you can modify the order of the displayed cases, and add or delete cases from the collection.

The acronym column in the **Multiple Cases** window displays the acronym associated with each case in the selection. You can assign an acronym to an AIMMS case through its **Properties** dialog box (see Section 17.2). Whenever available, AIMMS will use the acronym in multiple case displays within the graphical user interface.

The case acronym

17.5.1 Viewing multiple case data

The prime use of multiple case selection takes advantage of AIMMS' capability of displaying data from multiple cases within its graphical objects. Figure 17.7 illustrates a table which displays the contents of a single identifier for all the cases in the case selection shown in Figure 17.6.

Viewing multiple case data

You can turn a data object, in the graphical end-user interface, into a multiple case object by checking the multiple case property in the object-specific options in the object **Properties** dialog box. Figure 17.8 illustrates the object-specific **Properties** dialog box of a table object. As a result of enabling multiple case display, the object will be extended with one additional virtual dimension, the case index, which will be displayed in a standard way.

Creating multiple case objects

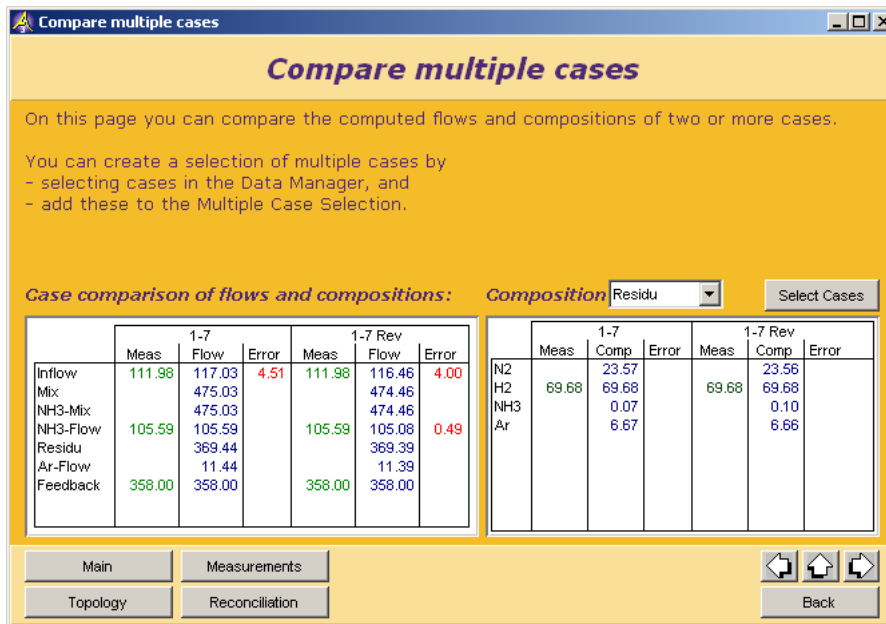


Figure 17.7: Example of a multiple case object

AIMMS only supports the display of multiple case data in object types for which the added dimension can be made visible in a well-defined manner. The most important object types that support multiple case displays are tables, pivot tables, curves, bar charts and scalar objects. Because of the extra dimension, the bar chart object is only able to display multiple case data for scalar and 1-dimensional identifiers. During a single case display, a bar chart can also be used to view 2-dimensional identifiers.

Restrictions

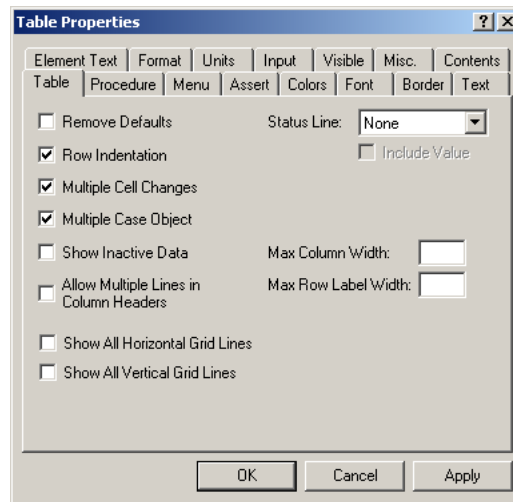
17.5.2 Case referencing in the language

In addition to viewing data from multiple cases as graphical objects in the graphical user interface, AIMMS also allows you to reference the data of cases that are not currently active within the model. This allows you, for instance, to perform advanced forms of case differencing by comparing the current values of particular identifiers in your model with the corresponding values stored in an inactive case.

Using inactive case data

The collection of all cases available in the **Data Manager**, is available in the AIMMS language through the predefined integer set `AllCases`. Each case in the **Data Manager** is represented by an integer element in this set, and, as explained in the Section 19.4.5, AIMMS offers several built-in functions to obtain additional information about a case through its case number.

The set AllCases

Figure 17.8: Table-specific **Properties** dialog box

You can reference the values of specific identifiers within a particular case by simply prefixing the identifier name with an index or element parameter in the set `AllCases`. Thus, if `cs` is an index in the set `AllCases`, the following simple assignment will inspect every case, and store the values of the variable `Transport(i,j)` stored in that case in the parameter `CaseTransport`, which has one additional dimension over the set of `AllCases`.

Referencing case data

```
CaseTransport(cs,i,j) := cs.Transport(i,j);
```

The capability of referencing inactive case data, enables you to perform advanced forms of case comparison, which would be hard to accomplish without the AIMMS facilities for case referencing. As an example, consider the following statement.

Advanced case comparison

```
RelativeDiff(cs,i,j) := (cs.Transport(i,j) - Transport(i,j)) /$ Transport(i,j);
```

It computes the relative difference between the current values of the variable `Transport(i,j)` and those values stored for each case on disk. You can display this data, for instance, in the graphical user interface.

AIMMS stores the case selection constructed in the **Multiple Case Selection** dialog box discussed in the previous section in the predefined set `CurrentCaseSelection`, which is a subset of the set `AllCases`. Thus, you can very easily apply the above possibilities to only the cases selected by your end-users in the **Multiple Cases** window. The following statement illustrates a small adaptation of the previous example to restrict the computation of the relative difference to only the cases in `CurrentCaseSelection`.

The set Current-CaseSelection

```
RelativeDiff(cs in CurrentCaseSelection,i,j) :=  
  (cs.Transport(i,j) - Transport(i,j)) /$ Transport(i,j);
```