
Introduction

He would keep on trying to do this or that with a grim persistence that was painful to watch . . .

John Wyndham, 'The Day of the Triffids'¹

Maple is a computer program capable of performing a wide variety of mathematical operations. It originated in the early 1980s as a computer algebra system, but today this description doesn't really do it justice. Maple has facilities for algebra, calculus, linear algebra, graphics (two- and three-dimensional plots, and animations), numerical calculations to arbitrary precision, and many other things besides. It is widely used in universities across the world, and is particularly useful for tasks that are tedious and error-prone when performed by humans, such as manipulating complicated series expansions and solving large sets of simultaneous equations. Used correctly, Maple can save time and quickly solve problems that would otherwise be intractable. Used incorrectly, it can lead to frustration, and the destruction of expensive IT equipment.

At the time of writing, the current version is Maple 2016. Versions before Maple 2015 were numbered starting from 1; the last of these was Maple 18. New features introduced in each version from Maple 4.0 onwards can be viewed using the help system (see Section 2.2). For the most part, recent changes have been relatively minor, at least as far as the material in this book is concerned. Consequently, all of the examples work with both Maple 2015 and Maple 2016. In fact, most will work in older versions as well, though naturally the number of exceptions increases the further back one goes. Two substantial new features are the `dataplot` command, discussed in Section 6.6, and the new rules concerning terminating characters, described in Appendix B (see also Section 2.3). Both of these were introduced in Maple 2015.

¹ Penguin Books, 1954. Reprinted by permission of Pollinger Limited (www.pollingerltd.com) on behalf of the Estate of John Wyndham.

1.1 Why This Book?

This book is intended for students, teachers and researchers who will ultimately wish to use Maple for advanced applications. Here, ‘advanced’ means something more complex than evaluating a single integral, but not necessarily designing and running a simulation of the latest jet engine. The book is suitable for undergraduates and postgraduates taking a course in which Maple is used, and for researchers who intend to use Maple for part of their work. It can also serve as a consolidation guide for users who already have some knowledge of Maple, but find themselves unable to decipher and eliminate certain error messages, or who currently rely on apparently magic recipes for solving problems, based on commands or operations whose meaning is not clear. There is no reliance on magic recipes here. Every feature we use is properly explained, with references to the online documentation where appropriate. The book is not a comprehensive reference guide (already available via the help system; see Section 2.2), nor is it a beginner’s guide in the normal sense. It most certainly is not a ‘guide for dummies’. We start from the beginning, assuming no prior knowledge of the subject whatsoever, but where advanced topics are central to understanding Maple they are tackled head-on, even as early as Chapter 2. In particular, the evaluation rules are a regular feature throughout. These determine the order in which input is processed by Maple. In most circumstances they are fairly simple (Section 8.4 discusses some more complex situations), but they are *absolutely crucial*.

Using this book, it is possible to tackle many complex problems without any additional Maple documentation. Readers can quickly progress to using packages and commands that have not been discussed, because the principles introduced apply across the whole system. Where the book alone is not sufficient, those who have read it will find themselves able to fully understand Maple’s help pages, which can be rather technical, and the Maple Programming Guide (see Section 2.2), which is squarely aimed at very advanced users.

1.2 The Maple Interface

Users may interact with Maple in several different ways. One may proceed by typing commands, by using interface driven methods based on palettes and context menus, or by a combination of the two. The advantages and disadvantages of a particular method are not always obvious, though they may be hugely significant.



Interface driven methods can be very tempting to new users, because they may appear to eliminate the need to study a manual. For example, with the factory setting for input (2-D Math; see Section 2.1) it doesn't take much effort to work out how to sum series and evaluate integrals, using the Calculus and Expression palettes at the sides of the screen (if these are not visible, they can be revealed by choosing **Palettes** ► **Expand Docks** from the **View** menu). However, most users will eventually want to try something a little more complicated, in which case things are not quite so straightforward. For all its power, Maple is only a computer program, and as such it can only understand mathematical input that is structured in the correct way. Just as it is possible to type incorrect commands, it is also possible to use the palettes incorrectly, and construct something that a human mathematician might understand, but Maple does not. One way or another, technical issues will sometimes arise, and a solid understanding of Maple is needed to deal with these effectively. In view of this, interface driven methods don't save much time, unless the software is to be used exclusively for solving elementary problems. There will be no further discussion of such methods here (search for the User Manual in the help system for more information on the subject).

This book takes a command driven, or programmatic, approach to Maple, with the focus on the language rather than the interface. This has two principal advantages. First, it scales up very easily: the simple building blocks that make up the Maple language can be assembled to solve complex problems in an efficient way. This is where the real power of Maple lies. Second, there is transparency: a Maple worksheet constructed using a sequence of mouse clicks and menu selections is opaque in that a user opening it cannot see immediately (if at all) how it was created, or how it could be modified and adapted to his/her needs.

On the other hand, a worksheet composed from typed commands is 100% transparent. It may be that some users who master the Maple language later decide that more interface driven methods are suitable for some or all of their work. However, such users will continue to find this book valuable, because understanding the Maple language makes the behaviour of its interface far more tractable.

1.3 How to Read This Book

The technical material in this book is intended to be read in order and in its entirety. Great effort has been expended to keep the content short, while still covering all of the key points. Time has also been spent minimising the number of situations in which concepts are used before they are properly introduced. In a few places, these structural aberrations turn out to be unavoidable (or the lesser of two evils). Where this is the case, the simplest possible examples have been used to illustrate the issue at hand, and a reference to a later section, in which the out of place concept is dealt with in detail, is always given. At the very least, every reader should study Chapter 2. Most of this is very basic, but many fundamental aspects of Maple are described here, and without knowledge of these its behaviour can seem mysterious at best, and infuriating at worst. Chapters 3–6 depend heavily on Chapter 2, but less so upon each other. The majority of users will need the material in Chapters 3 and 4, which introduce Maple’s symbolic computation facilities. Chapters 7–9 really need to be read in sequence. Without the ideas they contain, solving some problems will necessitate tedious, repetitive work, such as entering large numbers of very similar commands, which is not an efficient way to use Maple.

Throughout the book, items that appear in menus or dialogue boxes are shown in a rounded box with a grey background, such as . Keystrokes such as  have a sharp-cornered box with a white background. Icons for toolbar buttons are shown as they appear in Maple 2016; the Maple 2015 version is also shown or described if it is significantly different. Small blocks of text marked with the symbol ★ are tips. These are useful (often very useful) but not vital pieces of information

or advice. In some cases it is possible to deduce them from other parts of the book. Input and output is shown in the same style in which it appears in Maple itself (provided the configuration process explained in Section 2.1 is followed). In particular, Maple commands and statements are shown in a **typewriter typeface**. In a few places, Maple statements, or parts thereof, have been omitted in order to illustrate a larger or more general structure. Text in *italics* is used to give an indication as to what is missing. Output is omitted if showing it requires an inordinate amount of space. This convention is used extensively in Chapter 6, where plots drawn by Maple are not shown.

To get the most out of the examples throughout the book, it is necessary to execute them in Maple, and to experiment by modifying them. To save time for readers, the relevant files have been made available for download.² Typing the examples manually is fine, but it may be necessary to insert additional **restart** commands (see Section 2.11) between some of them to prevent unintended interactions, especially if the ordering is changed. The online worksheets already contain sufficient **restart** commands to ensure that everything works exactly as shown. Copying and pasting from an electronic version of the book may lead to unexpected results, and is therefore not recommended.

Please report errors to ian.thompson@liverpool.ac.uk.

² pcwww.liv.ac.uk/~itho17/understanding_maple