Spreadsheet structure inspection using low level access and visualisation

Daniel Ballinger, Robert Biddle, James Noble
School Mathematical and Computer Sciences
Victoria University of Wellington
PO Box 600, Wellington 6001, New Zealand
{db, robert, kjx}@mcs.vuw.ac.nz

Abstract

Spreadsheets are an extremely common form of end-user programming used for many applications from student marks to accounting for global multinationals. Ways of studying the structure of a spreadsheet itself is normally constrained to the tools provided in the spreadsheet software. We wanted to explore ways to use new visualisations for spreadsheets, and this paper documents our approach.

Keywords: spreadsheets, visualisation, end user programming

1 Introduction

When first interacting with a spreadsheet, a user can have to process a daunting amount of information in terms of layout and hidden inter-cell dependencies created by formulas. This problem of the organisation of code has been observed by users for some time and was commented on by Bonnie Nardi [9] where she mentions:

> *It is difficult to get a global sense of the structure of an individual formula that may have dependencies spread out all over the spreadsheet table. Users have to track down individual cell dependencies one by one, tacking back and fourth all over the spreadsheet.*

In this paper we present a solution to this problem through the creation of a set of images that represent the contents at a more abstract level than possible with the spreadsheet software. The user would start with very general information and then progress towards the actual details present in the spreadsheet, all the time learning about the layout and dependency structures without being exposed to actual values and other lower level structures in the spreadsheet. To create these abstract diagrams access to the internals of the spreadsheet, at the level accessible by the user, is required. To avoid the limitations present in the spreadsheet software this information will be extracted to a more versatile programming environment. Java was chosen in the case of this project, as the extraction process will eventually form an integral part of a larger Java based toolkit that will be used to perform analysis of spreadsheets and it is important that it integrate tightly.

Copyright © 2002, Australian Computer Society, Inc. This paper appeared at *Fourth Australasian User Interface Conference (AUIC2003)*, Adelaide, Australia. Conferences in Research and Practice in Information Technology, Vol. 18. Robert Biddle and Bruce Thomas, Eds. Reproduction for academic, not-for profit purposes permitted provided this text is included.

2 Accessing the internal structure of a spreadsheet

Excel files are formatted and stored using one of several versions of the BIFF (Binary Interchange File Format) depending on the version of Excel used to save the file. There has already been significant work on documenting this file format (outside Microsoft), including that done by Daniel Rentz as part of the OpenOffice project [12].

2.1 What needs to be extracted

The artefacts required from each Excel file are derived from the low-level structure of the spreadsheets. The basic unit of interest is the cell, the atomic building block of the spreadsheet model. Each occupied cell will have a value that is of some type, such as text or a number. Cells that are to perform a calculation, where the result is assigned to the value, will also contain a formula. The formula extracted should be in exactly the same notation presented to the user when interacting directly with Excel. That is, not the stored Reverse-Polish-Notation (RPN) format or missing any important elements such as absolute reference identifiers. This requirement stems form the need to see the formula as the user would, so the full set of dependencies can be extracted using the toolkit's formula parser.

It is undesirable to build a full extraction utility to read the basic file format from scratch, as that is a complex project in itself.

Due to the limited budget for this research project the API needed to be freeware or open source. Several commercial packages for reading Excel do exist [15].

When looking for a suitable method of converting BIFF encoded files to a representation using Java object structures several technologies and processes were investigated. Each of these brought certain implications, benefits and limitations. All techniques were assessed on the ability to extract the required information, minimal user interaction, platform independence and fault tolerance. Platform independence is important so the extraction process can migrate comfortably with the rest of the Java based toolkit. Fault tolerance is important because some extraction methods encounter difficulty with certain Excel structures and it is important that they tolerate them to allow the toolkit to continue functioning.

The two third party libraries that follow provide the toolkits current extraction abilities.

2.1.1 IBM alphaworks Java Bean

IBM alphaworks have produced the ExcelAccessor Java bean suite [5] to access and modify the contents of worksheets using a Windows DLL. This option was trailed first as it had easy configuration options over the other available methods. The bean usually works well, but the use of native code and requirement of having Excel installed limits its overall portability. Also, it is prone to irregular crashes during larger processing jobs involving multiple spreadsheets. When it does crash it often leaves an Excel process open on Windows 9x with subsequent runs having problems due to the still running process.

In its favour is the ability to correctly extract the exact components of the spreadsheet file that the user sees.

2.1.2 JExcelAPI

An alternative extraction process found uses JExcelApi by Andy Khan [8]. Being pure Java it integrates well with the rest of the toolkit and doesn't require native code. Some element types produce problems with the extraction methods currently available, such as array functions, intersections, absolute references that become relative in formulas, and files using a BIFF from early versions of Excel and the more recent XP version.

2.1.3 Future opportunities

The capability to extract values and string formulas in the user observed format from all occupied cells of the current range of freeware Java Excel readers is improving constantly as most are still actively in development. An example is the Apache POI project [10] where formula support is still being added.

3 Aspects of Spreadsheet structure

The spreadsheet paradigm exhibits two main characteristics: the spatial arrangement of cells on a table and the logical relationships between them created by formulas. These characteristics are not entirely disjoint, with the spatial relationship between cells often having a strong correspondence to the logical dependencies between them. A third, temporal, characteristic is exhibited by a subset of spreadsheets that are long-lived and contain a large number of cells with complex dependencies. Markus Clermont et al. observe that due to the nature of these long lived sheets they can go through similar evolutionary steps as conventional software [3].

Studies have investigated how spreadsheets assist usability.

The benefit of the spreadsheet paradigm is that it allows users to represent their understanding of the problem space on a two-dimensional table. This visual tabular layout is augmented via invisible formulas [6]. While entering raw data and formulas into the spreadsheet, users are shielded from the low-level details of traditional programming [1]. Application domain experts are able to express themselves in a reasonably transparent programming environment in a similar fashion to how they would with pen, paper, and a calculator. In doing this, the developer is naturally attempting to organise the arrangement of a sheet in a way related to its semantics [3].

4 Software support for visualisation with examples

The following figures are examples of images produced by the toolkit. For the purposes of demonstrating how the diagrams are derived from a single Excel workbook the source spreadsheet has been specially constructed using 3 worksheets of hypothetical data and dependencies.

4.1 Real-estate utilisation

When first encountering an unfamiliar workbook it can be useful to get a general orientating view of the layout of all the worksheets.

By counting the number of occupied cells that occur at each spreadsheet coordinate in each worksheet a *real-estate* diagram is created. The main purpose of this diagram is to understand basic positioning information about the workbook. Using this diagram it is possible to quickly get an impression about which sections of the spreadsheet are occupied and to what extent these sections are utilised. This diagram also makes finding sections of data that the designer has purposely tried to obscure easier to find. These obscured sections will often be dragged to a far corner of the sheet and provide functionality that does not need to be seen when using the core of the sheet, such as lookup tables.

In the following diagrams, the cells in the spreadsheet are arranged on the grid lines in their corresponding spatial positions. This produces a diagram with a layout very similar to that presented in Excel. The advantage of this is that relationships understood in one view could more easily be carried over to the view presented in Excel.
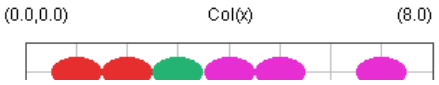
Two alternative methods are available to display this information.