VICTORIA UNIVERSITY OF WELLINGTON

Te Whare Wananga o te Upoko o te Ika a Maui

School of Mathematical and Computing Sciences

Computer Science

Spreadsheet structure inspection using low level access and visualisation

Daniel Ballinger and Robert Biddle and James Noble

Technical Report CS-TR-02/22
September 2002

School of Mathematical and Computing Sciences
Victoria University
PO Box 600, Wellington
New Zealand

Tel: +64 4 463 5341
Fax: +64 4 463 5045
Email: Tech.Reports@mcs.vuw.ac.nz
http://www.mcs.vuw.ac.nz/research

VICTORIA UNIVERSITY OF WELLINGTON
Te Whare Wananga o te Upoko o te Ika a Maui

School of Mathematical and Computing Sciences
Computer Science

PO Box 600
Wellington
New Zealand

Tel: +64 4 463 5341, Fax: +64 4 463 5045
Email: Tech.Reports@mcs.vuw.ac.nz
http://www.mcs.vuw.ac.nz/research

Spreadsheet structure inspection using low level access and visualisation

Daniel Ballinger and Robert Biddle and James Noble

Technical Report CS-TR-02/22
September 2002

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.

Publishing Information

This paper was included in the Proceedings of the 4th Australasian User Interface Conference, Adelaide, South Australia, February 2003; Conferences in Research and Practice in Information Technology, Vol 18.

Author Information

Daniel Ballinger was a BSc(Hons) student in Compupter Science at Victoria University of Wellington in 2002. Robert Biddle and James Noble are members or the academic staff in the School of Mathematical and Computing Sciences at Victoria University of Wellington.

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
Email: {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 (Nardi 1993) 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.

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 (Rentz n.d.).

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 (ACTUATE 2002).

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 (alphaworks ibm. n.d.) 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 (Khan n.d.). 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 (Oliver & Barozzi 2002) 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 (Clermont, Hanin & Mittermeir 2002).

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 (Igarashi, Mackinlay, Chang & Zellweger 1998). While entering raw data and formulas into the spreadsheet, users are shielded from the low-level details of traditional programming (Ayalew, Clermont & Mittermeir 2000). 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 (Clermont et al. 2002).

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.

The first approach makes use of Java's standard graphics environment to draw grid-based 2D data as in figure 1.

In this diagram the number of cells that occur at each grid coordinate are counted and displayed using a heat map colour. Grid coordinates where a large number of cells occur will be coloured towards the red end of the colour spectrum will those with lower counts will be coloured towards the violet/blue end. Each row and column in this figure has a 1 to 1 relationship with the rows and columns in the spreadsheet respectively.

The second real-estate diagram utilises a 3D surface map approach created using Java3D (Sun 2002) and VisAD (Hibbard 2002), a visualisation tool for numerical data. This has the benefit of being interactive by allowing the user to rotate and zoom the image. An example of this type of diagram is figure 2. In this figure the left axis contains the rows and the upper (obscured) axis the columns. The altitude represents the occupancy level and is coloured to create a terrain type appearance.

4.2 Formulas at a Sheet Level

Once the user understands which parts of the spreadsheet are occupied the formulas that occupy each sheet can be examined.

4.2.1 Referencing cells

Each dependency created by logical references between cells in formulas can be converted to a mathematical vector. The vector between cells has an origin

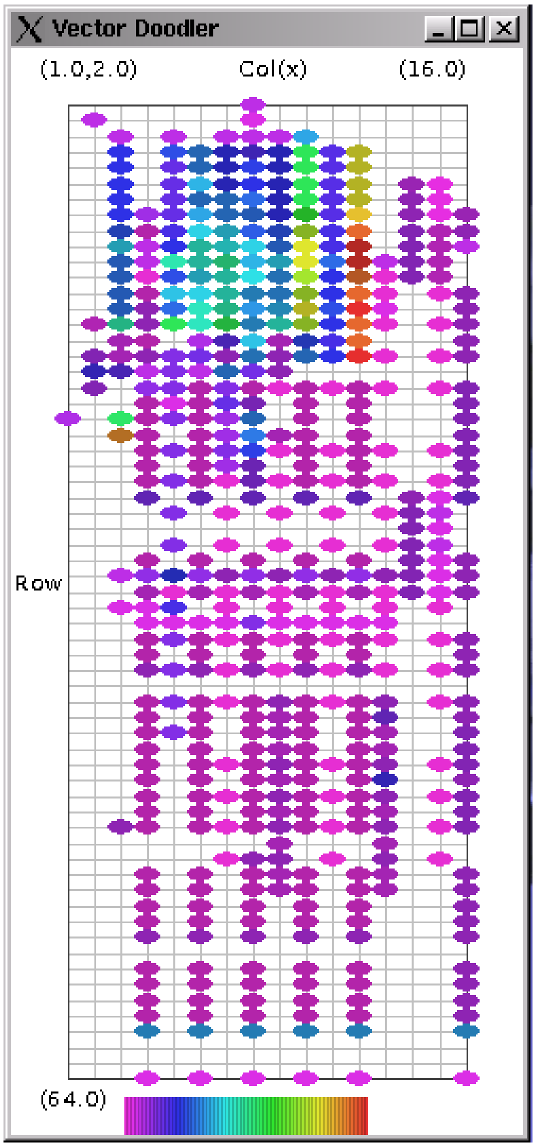

Figure 1: Real-estate utilisation diagram in 2D

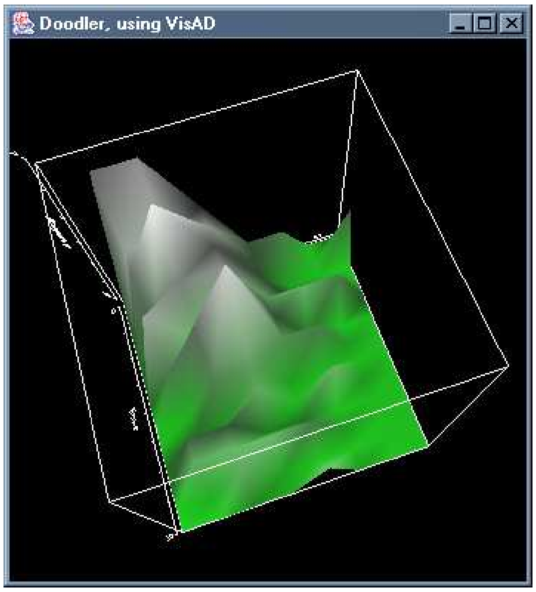

Figure 2: Real-estate utilisation diagram in 3D

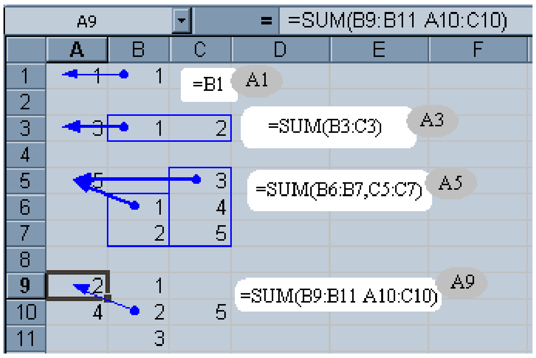

Figure 3: Excel's precedent trace auditing tool for sheet 1.

set at the referencing cell and a magnitude that ends the vector to the referenced cell.

Figure 3 shows the first sheet in the example workbook with precedent trace arrows added by Excel's built in auditing tools. We have added boxes by hand containing the formula present in the space immediately to the right of each cell. This particular sheet shows the 4 main techniques that a user can embed in a formula to reference other cells. Notice how Excel's trace is somewhat deceptive in the case of an intersection, in this case looking more like an individual cell reference.

In figure 4, created by the toolkit, a precedent cell dependency is represented by a blue arrow in a similar fashion to that present in Excel. Each cell containing a formula has a circle that represents the length of the formula as a string and hence the intersection has a larger circle than the union. Any ranged references are shown using a shaded box. A single range is shaded light grey while a union has the left and right sub-ranges coloured blue and green respectively. Intersections use yellow and red boxes for the left and right sub-ranges. The actual resulting intersection is shaded dark blue.

4.2.2 Absolute and relative references

Figures 5 and 6 demonstrate the colouring of relative and absolute cell references. The reference from B1, coloured blue, is a standard relative reference. B3 has a red absolute reference. A2 and C2 are both partially absolute references, in that only one axis is fixed with the $ symbol, and are coloured yellow and green respectively.

The distinctive colouring of these different styles of referencing should help in identifying patterns on large more complex sheets.

4.2.3 A more realistic Summation example

Figures 7 and 8 demonstrate a more realistic style of worksheet where a series of columns are summed and then crosschecked with the sum of the rows. This style of image is similar to the static global view by Takeo Igarashi et al. (Igarashi et al. 1998).

4.3 Related Work

The Holy Grail in spreadsheet research at the moment appears to be the admirable goal of detecting and correcting errors (Ayalew et al. 2000, Rothermel, Cook, Burnett, Schonfeld, Green & Rothermel 2000, Clermont et al. 2002, Igarashi et al. 1998, Nardi 1993,

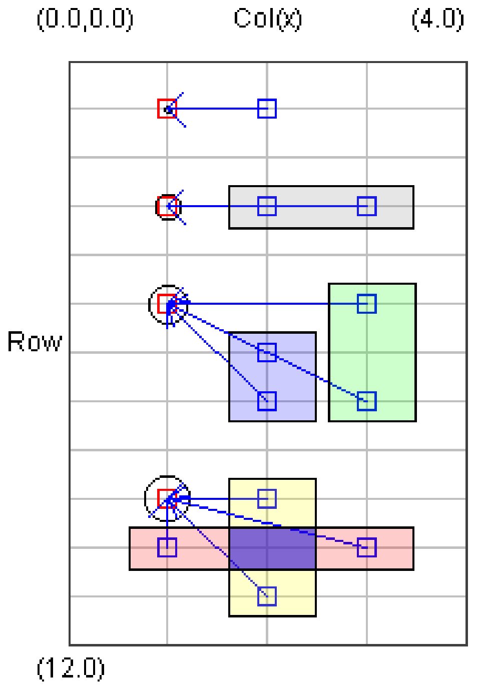

Figure 4: Formula Extracted from Sheet 1.

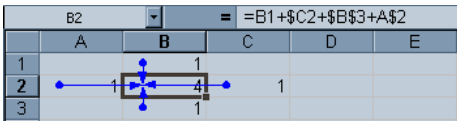

Figure 5: Sheet 2 with precedents trace.

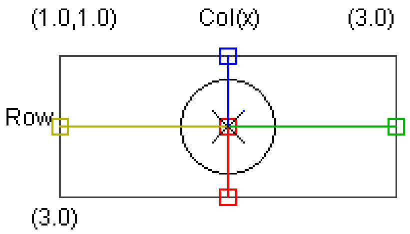

Figure 6: Formula Extracted from Sheet 2.

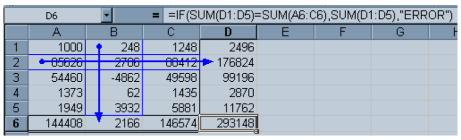

Figure 7: Sheet 4 with precedents trace.

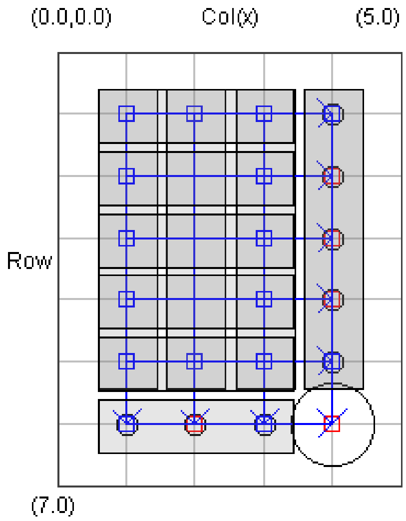

Figure 8: Formula Extracted from Sheet 4.

Panko 1998).

Margaret Burnett is continuing the development of a visual testing methodology to complement the visual programming environment that spreadsheets use (Rothermel et al. 2000). The motivation for this comes from the high degree of errors that are often present in spreadsheets.

Takeo Igarashi presented a series of techniques to address the hidden dataflow graphs and superficial tabular layouts of spreadsheets in (Igarashi et al. 1998). They were designed to improve the users understanding of the dataflow structure by enabling them to visually interact with the obscured structures. He mentions that there is a need to integrate these diagrams into a more realistic spreadsheet program. We have addressed this by allowing the diagrams to be generated free from the application.

Jorma Sajaniemi has presented a theoretical model of spreadsheets along with a description of various spreadsheet auditing mechanisms employing the model (Sajaniemi 2000). Our toolkit could be used to support such mechanisms.

5 Conclusions

This paper details the creation of a tool to first extract artefacts from spreadsheets and then convert this basic information into visualisations. Its purpose is to augment that provided already by Excel and make it easier to then create new diagrams.

The extraction methods and basic visualisations presented in this paper are encompassed in a larger suit of tools intended to be used for corpus based visualisations.

References

ACTUATE (2002), 'Formula one e.spreadsheet engine'. `http://www.actuate.com/f1/products/espreadengine.asp`.

alphaworks ibm. (n.d.), 'Excelaccessor'. `http://www.alphaworks.ibm.com/ab.nsf/bean/ExcelAccessor`.

Ayalew, Y., Clermont, M. & Mittermeir, R. T. (2000), Detecting errors in spreadsheets, *in* 'EuSpRIG

2000 Symposium: Spreadsheet Risks, Audit and Development Methods'. University of Greenwich, London. http://citeseer.nj.nec.com/485147.html.

Clermont, M., Hanin, C. & Mittermeir, R. (2002), A spreadsheet auditing tool evaluated in an industrial context, *in* 'EuSpRIG 2002 symposium'. http://www.sysmod.com/eusprig02.htm.

Hibbard, B. (2002), 'Visad, java component library'. Space Science and Engineering Center - University of Wisconsin - Madison. http://www.ssec.wisc.edu/~billh/visad.html.

Igarashi, T., Mackinlay, J. D., Chang, B.-W. & Zellweger, P. (1998), Fluid visualization for spreadsheet structures, *in* 'Visual Languages', pp. 118–125. http://citeseer.nj.nec.com/igarashi98fluid.html.

Khan, A. (n.d.), 'Excelread - a java api to read excel 97 spreadsheets'. http://www.andykhan.com/excelread/index.html.

Nardi, B. A. (1993), *A Small Matter of Programming: Perspectives on End User Computing*, MIT Press. http://www.darrouzet-nardi.net/bonnie/ASmallMatter.html.

Oliver, A. C. & Barozzi, N. K. (2002), 'Hssf (horrible spreadsheet format) part of poi'. Apache Software Foundation. http://jakarta.apache.org/poi/hssf/index.html.

Panko, R. R. (1998), What we know about spreadsheet errors, *in* 'Journal of End User Computing.'. http://panko.cba.hawaii.edu/ssr/Mypapers/whatknow.htm.

Rentz, D. (n.d.), 'Openoffice.org's documentation of the microsoft(r) excel file format'. http://sc.openoffice.org/.

Rothermel, K. J., Cook, C. R., Burnett, M. M., Schonfeld, J., Green, T. R. G. & Rothermel, G. (2000), WYSIWYT testing in the spreadsheet paradigm: an empirical evaluation, *in* 'International Conference on Software Engineering', pp. 230–239. citeseer.nj.nec.com/rothermel99wysiwyt.html.

Sajaniemi, J. (2000), 'Modeling spreadsheet audit: A rigorous approach to automatic visualization', *Journal of Visual Languages and Computing* **11**(1), 49–82. citeseer.nj.nec.com/sajaniemi98modeling.html.

Sun (2002), 'Java 3d(tm) api'. Sun Microsystems, Inc. http://java.sun.com/products/java-media/3D/.