

1 Introduction

A key question in human-computer interaction research is: How can we design software so that non-programmers can build their own applications? One way to study this question is to analyze existing user programming environments such as spreadsheets, HyperCard, the Metaphor Capsule, New Wave, fourth generation database management systems, style sheets in word processing programs, and statistical packages such as SPSS and SAS. We can learn much from their successes and failures. Unlike most research systems, these commercially available programs have large user populations, enabling us to study the programs as they are actually used in offices and homes.

Our work focuses on spreadsheet software. Spreadsheets have proven enormously popular with personal computer users, and their benefits have been enumerated by Kay (1984), Hutchins, Hollan and Norman (1986), and Lewis and Olson (1987). These investigators have noted that spreadsheets provide a concrete, visible representation of data values, immediate feedback to the user, and powerful features such as applying formulas to blocks of cells.

These characteristics of spreadsheets are important. However, the single biggest advantage of spreadsheets is not cognitive but motivational: after only a few hours of work, spreadsheet users are rewarded by simple but functioning programs that model their problems of interest. Many users lack formal programming education, and perhaps more importantly, they lack an intrinsic interest in computers. The key to understanding non-programmers' interaction with computers is to recognize that non-programmers are not simply under-skilled programmers who need assistance learning the complexities of programming. Rather, *they are not programmers at all*. They are business professionals or scientists or other kinds of domain specialists whose jobs involve computational tasks. It is not enough to say that these users need systems that are "easy to use." User programming systems should allow users to solve simple problems within their domain of interest in a few hours.

2 Methodology

To understand user programming, we believe it is necessary to find out how people actually use software in the everyday contexts of homes and offices. We have chosen to study a small number of people in some depth to learn how they construct, debug, and use spreadsheets. We are interested in the kinds of problems for which people use spreadsheets and how they themselves structure the problem solving process – topics that by their very nature cannot be studied under the controlled conditions of the laboratory.

For the study, we interviewed and tape recorded conversations with spreadsheet users in their offices and homes, and collected examples of their spreadsheets.¹ Study participants were found through an informal process of referral. We told prospective participants that we are interested in software for non-programmers and that we want to talk to people actively using spreadsheets. The interviews were conversational in style, intended to capture the users' experiences in their own words. A fixed set of open-ended questions was asked of each user, though the questions were asked as they arose naturally in the context of the conversation. Part of the interview session was devoted to viewing users' spreadsheets on-line and discussing their uses and construction. Material in this paper is based on about 335 pages of transcribed interview material from a total of eleven users. User names given

in this paper are fictitious.

3 Why do spreadsheets work so well?

The usefulness of spreadsheets derives from two properties of their design:

- Computational techniques that match users' tasks and that shield users from the low-level details of traditional programming, and
- A table-oriented interface that serves as a model for users' applications.

The power of spreadsheets comes from the combination of these properties – either separately would be inadequate to solve the spreadsheet user's two basic problems: *computation* and *presentation*. Our intent in this paper is to understand the reasons for the success of spreadsheets, and to look for general principles that might be applied to other user programming environments.

4 Computation: The spreadsheet formula language²

The spreadsheet formula language allows users to compute values in their models by expressing relations among cell values.³ To use the formula language, the user must master only two concepts: cells as variables, and functions as relations between variables. With relatively little study, the user acquires the means to solve the basic computational problems of any modeling task: creating entities that represent the variables in the problem, and expressing relations among the entities.

As users become more proficient at using spreadsheets, they learn more advanced programming concepts such as relative and absolute cell references, iteration, and conditionals. Knowing these more advanced concepts is immensely useful, but we want to emphasize that they are not necessary for beginning users building simple spreadsheets. Spreadsheets allow users to perform useful work with a small investment of time, and then to go on to more advanced levels of understanding as they are ready. In our research we found that users may add new programming concepts to their repertoire very slowly.

For example, Jennifer, a user in our study, has been using spreadsheets for about five years. She has an accounting position of considerable responsibility in a high technology firm. Jennifer knows how to write nested conditionals, to link individual spreadsheets, and to create simple macros. She has not, however, learned how to iterate operations over a cell range (a group of contiguous cells), though she is aware of this capability, and plans to learn how to do it. Despite what would be a fatal gap in her knowledge in a traditional programming language, Jennifer is a successful spreadsheet user. Moreover, she is continuing to expand her knowledge about spreadsheets, at her own pace.

The spreadsheet formula language is characterized by:

- High-level, task-specific functions.
- Very simple control constructs.

This combination of attributes in the spreadsheet formula language strikes a fine balance between expressivity and simplicity. Users have sufficient means to model their problems, but at a very attractive price in terms of learning and development time.

4.1 High-level, task-specific functions

The formula language offers a small number of arithmetic, financial, statistical, and logical functions. Most spreadsheets also offer simple database functions, date and time functions, and error trapping functions. In our study we found that most users normally use fewer than ten functions in their formulas. Users employ those functions pertinent to their domain (e.g., financial analysis) and do not have need for other functions.

Spreadsheet users are productive with a small number of functions because the functions are high-level, task-specific operations that do not have to be built up from lower level primitives. For example, a common spreadsheet operation is to sum the values of a range of cells within a column. The user writes a simple formula that specifies the sum operation and the cells that contain the values to be summed. The cell range is specified compactly by its first and last cell; e.g., `SUM(C1:C8)` sums cells 1 – 8 in Column C. In a traditional programming language, computing this sum would require at least writing a loop iterating through elements of an array, and creating variable names for the loop counter and summation variable. Spreadsheet functions obviate the need to create variable names (cells are named by their position in the grid), and to create intermediate variables to hold results – non-task-related actions that many users find confusing and tiresome (Lewis and Olson, 1987).

4.2 Control constructs

Lewis and Olson (1987) stated that a strength of spreadsheets is the “absence of control model.” They noted that flow of control is a difficult programming concept. However, it is not the case that control mechanisms are absent in spreadsheets – they are just conceptually simple. Formulas can be written as if-then-else statements. A conditional in a spreadsheet formula is easy to understand because it does not transfer control from one part of the spreadsheet to another; its effects are local to the individual cell. Iteration is also quite simple; users can select a range of cells over which to iterate an operation. These conditional and iterative capabilities were routinely used by those in our study and added a great deal of functionality to their spreadsheets.

Users’ own formulas contain an implicit flow of control as any arbitrary cell can be related to any other. When a value in one cell changes it may trigger a series of changes in dependent cells. This is the very basis of the spreadsheet’s functionality, and it is quite powerful. Users can begin by building quite simple cell relations, and move on to more elaborate models as their knowledge expands.

4.3 User Programming Languages

Several lines of research on user programming proceed from the assumption that graphical techniques such as as program visualization, visual programming (Myers, 1986), and program induction (Maulsby and Witten, 1989)⁴ will provide significant leverage to non-programmers. Since the spreadsheet formula language is textual, we question whether graphics *per se* is really the key to user programming languages. A limited language of high-level functions is more important than the particular form the language takes. In our study we asked users to discuss the disadvantages of spreadsheets. Not one user mentioned difficulties with the formula language (though users had other specific complaints

about spreadsheets). Syntax is often suggested as a problem area in textual languages, but in our study users reported that syntax errors were few once they were familiar with a spreadsheet. Users noted that in any case, most such errors are immediately caught by the spreadsheet itself which will not permit poorly formed formulas. Proper syntax checking appears to be sufficient to enable users to cope with syntax errors. Textual languages are compact, efficient, and can be developed in less time than graphical languages. These are significant advantages which should be considered in the development of user programming languages.

Another line of user programming research suggests that users can program by modifying existing example programs (Lewis and Olson, 1987; Neal, 1989). These researchers believe that the complexity and difficulty of general purpose programming can be reduced by giving non-programmers a “head-start” with existing code which they then modify for their particular applications.⁵ Even assuming that the daunting problem of information access were to be solved such that users could easily locate apposite examples, we think that general purpose programming languages, no matter how well supported, are not appropriate for the large population of users who lack intrinsic interest in computers, and have very specific jobs to accomplish. These users should be supported at their level of interest, which is to perform specific computational tasks, not to become computer programmers.

What about programming by modifying domain-specific examples? This solution still does not solve the problem of having to depend on the existence of appropriate example code. True user programming systems allow users to build a meaningful application without reliance on obtaining code from other more sophisticated users. No programmer wants to lack the skills with which to create a program from scratch, since that is so often necessary. If it is impossible to begin a program without an existing program, the user is denied real control over the computational environment. It is not clear whether users who modify existing example programs ever really come to understand the programs they modify. Without a firm grasp of the language in which the examples are written, the ability to modify a program to suit one’s needs would seem very limited.

There is a need to draw a distinction between programming by modifying example programs and the reuse of software modules. Reusable software modules are clearly desirable. In the spreadsheet world there is software reuse in the form of templates used by groups of users (Nardi and Miller, 1989). However users are not dependent on templates, and they routinely create their own applications using the spreadsheet formula language. We asked users what they liked about spreadsheets, and several users reported that they can be “creative” with spreadsheets, that it is “easy” to build their own models. One user captured a general feeling about spreadsheets in noting that he thinks of the spreadsheet as a “blank canvas” – a medium in which to directly express his own thoughts; just the opposite of an artifact created by someone else that must be re-worked before it is of any use.

The large variety of applications modeled with spreadsheets (Lewis and Olson, 1987) does indeed suggest a blank canvas. Spreadsheet applications include mathematical modeling (Arganbright, 1986), simple databases, managing small businesses, forecasting trends (Janowski, 1987), analyzing scientific and engineering data, and of course the financial applications for which they were first intended. Users have programmed these diverse and sometimes sophisticated applications (see Arganbright, 1986) without the aid of example programs.

5 Presentation: The tabular grid

The second major element of the spreadsheet interface is a strong visual format for organizing and presenting data – the tabular grid into which users put data values, labels and annotations. The table helps users solve three crucial problems: viewing, structuring and displaying data.

5.1 Viewing data

Virtually every user in our study reported that an advantage of spreadsheets is the ability to view large quantities of data on one screen. Applications modeled with spreadsheets are data-rich, and users in our study had a strong preference for being able to view and access as much data as possible without scrolling the screen. How do spreadsheets manage large amounts of data such that users feel that it is well-presented and comprehensible?

Spreadsheets have done well at data display by borrowing a commonly used display format – that of the table. Cameron (1989) pointed out that tables have been in use for 5000 years. Inventory tables, multiplication tables and tables of reciprocal values have been found by archaeologists excavating Middle Eastern cultures. Ptolemy, Copernicus, Kepler, Euler, and Gauss used tables. Modern times brought us VisiCalc, the first personal computer spreadsheet. VisiCalc was modeled directly on the tabular grid of accountants' columnar paper which contains numbered rows and columns. It is interesting that today's spreadsheets, while much enhanced in functionality, have not changed the basic VisiCalc format in the smallest detail. A tabular grid in which rows are labeled with numbers and columns are labeled with letters characterizes all commercially available spreadsheets.

Tables are so familiar and common in our everyday lives that we are unlikely to pause to appreciate their clever design – they are extraordinarily simple and viewable. It is quite easy, even in a large table, to ascertain the categories represented in the vertical and horizontal dimensions, to scan for individual data values, and to get a sense of the range of values and other characteristics such as a rough average. The perceptual reasons for tables' exceptional capability to effectively display data are not well understood, but Cleveland's notion of "clustering" – the ability to hold a collection of objects in short-term memory and carry out further visual and mental processing (personal communication, 1989) – seems relevant. The arrangement of data items in rows and columns appears to permit efficient clustering, as users can remember the values in a row or column and then perform other cognitive tasks that involve the values.

The familiarity of tables further enhances the ease with which we use them; our schooling explicitly trains us in table use from reading calendars to learning matrix algebra, and everyday experience provides ample opportunity to both create and view tables.

Tables provide good information access as users can locate data in a simple geometric space. In a large spreadsheet though the data are not continually visible (a desideratum of proponents of direct manipulation) as the entire spreadsheet will not fit on one screen, the geometric organization of the grid permits users to find their data quickly.

For example, Jennifer was discussing a spreadsheet that contained about 300 rows and we asked her how she "gets around" in this large spreadsheet. Notice that in the following exchange she thinks we want to know the mechanics of navigating with mouse and keyboard.

She adds the comment about the geometric layout of her spreadsheet as a clarification, though we have not talked about layout at any time in her interviews:

Interviewer: Now when you're actually using a spreadsheet this big, how do you get around to the places you want to be?

Jennifer: I use the mouse on the gray bar. It lets you leap down a page. It's kind of like page-up and page-down. But I can page-right and -left more easily than you can with the keyboard.

Interviewer: OK, so that's not really an issue. Even though you do have a lot of data it's pretty easy to find it.

Jennifer: UmmHmm. I'm so familiar with the spreadsheet too, that I know that if I'm here [points to a place on the spreadsheet] in Municipal Bonds, that I know I'm in the middle of the document, and I know that Preferred Stocks is above that, and I know that Collateralized Mortgage Obligations are below that, so depending on what the next transaction is, I know whether to go up or down.

Today's spreadsheets also allow users to assign names to cell ranges. In our study, some users assigned range names in large spreadsheets and then located the range by typing in the name, rather than scrolling to it. Spreadsheets thus offer both logical and spatial clues to data location that enable users to quickly find data even in very large spreadsheets.⁶

5.2 Structuring data

A spreadsheet table is much more than an effective data display – it is a problem solving medium. We tend to think of data presentation as largely a matter of setting forth information for the user to view and browse. But the means by which data are presented strongly affect the problem solving process. The table inhabits problem solving cognition in two ways.

First, the very structure of the table is the means by which users come to organize their models. Data are arranged into rows, columns, and cells. The spreadsheet provides a structure into which a model is cast. *Users do not have to invent a structure* – it is given to them. The initial phase of a modeling problem is reduced to simply recognizing a format into which a problem is framed, rather than being faced with the necessity of inventing a format from scratch.

Second, in the very process of laying out data in a spreadsheet, the user is viewing and studying the data which are immediately reflected back from the table. This contrasts sharply with traditional programming where the parameters and variables of a model are implicit in the procedures that manipulate them, and have no explicit visual representation.⁷

The structured visual format for data presentation provided by the spreadsheet table plays an active role in helping users to *structure* and then *critique* their models. We now look in more detail at how users structure and critique spreadsheet models.

The most striking thing about spreadsheets is how they help users to think through problems. *The tabular format provides a simple but powerful framework onto which users map their problems.* The importance of this structure became evident as users described how they use spreadsheets to model problems *even when their initial ideas about the problem*

solution are extremely ill-defined. We were struck by the fact that many users reported that when they begin a spreadsheet they have only a general goal in mind (e.g., “maximize profits over the next three quarters” or “decide how much house we can afford”), and very little idea of how to achieve the goal. When beginning work on a new spreadsheet, users often do not even know what the parameters of a problem are. They only find out about all the relevant aspects of a problem in the process of actually trying to solve it.

For example, Jeremy, one of the participants in our study, described how he learned to use spreadsheets. A job assignment required developing business plans for joint ventures with foreign companies. Large, complex spreadsheet models were part of the plans. Although programmers were available to help, Jeremy discovered that not only was it easier to be in control of spreadsheet development himself, but that he could use the spreadsheet to work through the problem, in particular to identify the variables of interest and to make sure that the model was complete.

Jeremy described this process:

Jeremy: We had to have rather large complex spreadsheets [for the business plans] where you had lots of variables. And I found it easier to develop that myself than to go to somebody and say here’s what I want, here’s what I want, here’s what I want. And that’s what really got me going on [spreadsheets]. . .

Interviewer: Why was it easier for you to do this yourself than to specify it for a programmer?

Jeremy: . . . I think it was quicker and easier because I felt that I was learning as I went, as I was developing the spreadsheets, I was *learning* about all the variables that I needed to think about. It was [as] much a prop for myself as [a way of] . . . getting the outcome . . . And there were a lot of false *endings*, I should say, not false starts. I’d get to the end and think, “I’m done,” and I’d look at it and I’d say, “No, I’m not, because I’ve forgotten about one thing or the other.”

How do spreadsheets help users give shape to fuzzy ideas?

The spreadsheet provides an overall organizing framework of rows, columns and cells within which users organize the parameters, variables, formulas and subparts of their models.

Rows and columns are used to represent the main parameters of a problem. Users know that related things go in rows and columns, and all spreadsheet applications take advantage of the simple but powerful semantics provided by the row/column convention.

Each cell represents and displays one variable. A cell value may be a constant, or may be a calculated value derived from a formula. In the case of calculated values, the spreadsheet associates a visual object, the cell itself, with a small program, the formula. Program code is thus distributed over a visual grid, providing a system of compact, comprehensible, easily located program modules. The spreadsheet itself automatically updates dependent values as independent values change; therefore the user’s the task is to write a series of small formulas, each associated with a distinct visual object, rather than the more difficult task of specifying the full control loop of a program as a set of procedures.

Tables provide a simple mechanism for segmenting models into smaller subparts: leaving empty cells between segments. A spreadsheet can be modularized, at least visually, for the purpose of showing its subparts. In our study, we found that users segment spreadsheets by such criteria as years, months, geographic regions, companies, and departments.

Spreadsheets relieve users of the necessity of inventing their own modeling frameworks – a demanding task which would force them to build a problem solving infrastructure before getting to work on their actual goals. The spreadsheet table, by virtue of a structured visual format for presenting data, provides the hooks upon which a user hangs a model. The advantages of this structure are amplified by the fact that as the user builds the model, it emerges in a highly visible way. The model is not buried in a text file of many lines of computer code, it is not littered with obscure variable names, but instead consists of an orderly set of parameter names and variable values laid out in a simple two-dimensional space. Users can see exactly what their parameters and variables are as they add them to the model.

The visibility of the emerging model is very important in the problem solving process. In our study we found that users critique their models by visually inspecting them. As Jeremy noted, he evaluated the completeness of his model by *looking at it*: “I’d look at it and I’d say, ‘No, I’m not [done], because I’ve forgotten about one thing or the other.’” Other users also described the process of visually inspecting their models as they were building them. For example, one user stated: “I may not even . . . know the final form, look and feel of the spreadsheet that I want. I’ll just start getting the data in, and then I’ll start . . . playing with moving rows and columns around and doing things until *I see*, until I get what I want” (our emphasis). The act of viewing data in a spreadsheet table is thus not merely a means by which to find a data value, or check out the bottom line; it is a key aspect of the active process of model construction.

To summarize, the way data are presented – that is, what the user *sees* – shapes the problem solving process. As a user begins developing a spreadsheet, the tabular grid provides an overarching structure into which the parameters and variables of a model are cast. As the spreadsheet begins to take shape, the user views the emerging model and evaluates its accuracy and completeness. Within the framework of the rows and columns the user can restructure the model by re-arranging rows and columns and by adding new parameters as they become known. A spreadsheet model is grounded in the distinct tabular format of rows and columns, and is constructed in successive approximations as the user critiques the emerging model.

5.3 Displaying data

Users welcome structure in modeling the parameters and variables of their problems, but seek flexibility in creating their own displays. In office environments many spreadsheets are viewed and used by a group of co-workers, and users in our study emphasized the importance of creating effective presentations – often paper copies or slides of their spreadsheets.

The spreadsheet table has some useful flexibility. All spreadsheets allow users to vary column width, and modern spreadsheets allow users to vary individual column widths and row heights. Spreadsheets allow users to split the screen so that non-contiguous portions of a spreadsheet may be viewed on the screen (or printed out) at once. Spreadsheet cells are flexible building blocks; they are used not only to hold variables, but also to display labels and annotations, and to segment large spreadsheets into subparts by means of empty cells, as noted. All users in our study used some or all of these capabilities. All users in our study who had spreadsheet products that give users control over color, fonts, shading and outlining used these techniques to highlight important data. In short, spreadsheets give users a reasonably good user interface toolkit.

6 Summary

Of course, spreadsheets are not without their problems, many of which derive from the same properties that give them their strength. The ability to build spreadsheets through assigning small pieces of code to specific cells means that it is difficult to get a global sense of the structure of the spreadsheet, which requires tracing the dependencies among the cells. Many users in our study described awkward pencil and paper procedures for tracing cell dependencies in debugging spreadsheets. For the same reason, spreadsheets are not particularly modular: since the code that implements a particular piece of a spreadsheet is distributed over a potentially large and unpredictable set of cells, it is difficult to reuse a piece of one spreadsheet in another new spreadsheet.

Nevertheless, the strengths of spreadsheets are profound. Spreadsheets suggest that two key characteristics for user programming environments are:

1. A limited set of carefully chosen, high-level, task-specific operations that are sufficient for building applications within a restricted domain, and
2. A strong visual format for structuring and presenting data.

The ability to create applications with only a few functions is an important benefit of spreadsheets. Users have specific tasks to accomplish within their domain of interest. They want functionality that matches those tasks at a high level such that they do not have to either learn or use lower level primitives. Task-specific functions allow users to develop quick facility with a program and to build a real application, however simple, in a short time. The motivational barrier is thus breached as users achieve rapid success. As users continue to use a program they are not constantly faced with the job of stringing together lower level functions as they work, but can concentrate on the actual problem solving itself.

Spreadsheets succeed because they combine an expressive high level programming language with a powerful visual format to organize and display data. The user is actively engaged with the spreadsheet table as a problem solving device throughout the process of model building. The tabular structure of rows, columns and cells provides a modeling framework. The visibility of the emerging model allows the user to monitor and evaluate its accuracy and completeness.

The spreadsheet experience suggests that general programming languages are not the answer for non-programmers. Users who lack intrinsic interest in computers and who have specific tasks to get done are more likely to respond to a software system that provides high-level functionality in their area of expertise than to tolerate the slow detour of a general programming language. Insofar as techniques such as program visualization, visual programming, programming by example modification and program induction support general programming, they are unlikely, in our view, to succeed at helping non-programmers gain increased computational power. A more fruitful line of endeavor is to identify ways to support the development of high-level, task-specific languages and appropriate visual formats for new user programming environments.

7 Notes

1. The interviews were conducted by the first author. We use the plural “we” here for expository ease.
2. We refer to “the formula language” because most spreadsheet programs have nearly identical languages which differ only in small syntactic details.
3. Spreadsheets also have macro languages, but they are used by many fewer users and do not constitute the basic interface to spreadsheet functionality that we are concerned with here.
4. Programs are “induced” by generalizing from concrete examples created by the user via graphical direct manipulation techniques.
5. Example modification is not the same as the use of didactic examples which have an important role in learning and enhancing skill in programming (and many other areas of endeavor). Didactic examples are especially helpful in learning language syntax as part of a larger program of study in which the fundamental concepts of a programming language are learned.
6. In our study the largest spreadsheets had about a thousand rows.
7. The spreadsheet provides immediate feedback. When the user changes data values, other values related through formulas are immediately updated. This compression of the test-evaluate-debug cycle is an important feature of spreadsheets, but one that has been discussed by other investigators (Hutchins et al., 1986; Lewis and Olson, 1987) so we do not expand on it here.

8 Acknowledgments

We are grateful for comments from Lucy Berlin, Martin Griss, Jeff Johnson, Nancy Kendzierski, Jasmina Pavlin and Craig Zarner.

9 References

- Arganbright, D. (1986). Mathematical modeling with spreadsheets. *Abacus* 3:4:18-31.
- Cameron, J. (1989). A cognitive model for tabular editing. OSU-CISRC Research Report, June, 1989. Ohio State University.
- Cleveland, W. (1989). Personal communication.
- Hutchins, E., Hollan, J. and Norman, D. (1986). Direct manipulation interfaces. In *User Centered System Design* (Eds. D. Norman and S. Draper). Erlbaum Publishers: Hillsdale, NJ.
- Janowski, R. (1987). Spreadsheets: An initial investigation. Internal Technical Report, Hewlett-Packard Laboratories, Bristol, England.

Kay, A. (1984). Computer software. *Scientific American* 5:3:53-59.

Lewis, C. and Olson, G. (1987). Can principles of cognition lower the barriers to programming? In *Empirical Studies of Programmers: Second Workshop* (Eds. G. Olson, S. Sheppard and E. Soloway). Ablex Publishing Corporation: Norwood, NJ.

Maulsby, D. and Witten, I. (1989). Inducing programs in a direct-manipulation environment. In *Proceedings of CHI'89, Conference on Human Factors in Computing Systems*. April 30 - May 4, 1989. Austin, Texas. Pp. 57-62.

Myers, B. (1986). Visual programming, programming by example, and program visualization: A taxonomy. In *Proceedings of CHI'86 Conference on Human Factors in Computing Systems*. April 13 - 17, Boston.

Nardi, B. and Miller, J. (1989). Twinkling lights and nested loops: Distributed problem solving and spreadsheet development. Hewlett-Packard Laboratories, Palo Alto, STL-Report 89-30.

Neal, L. (1989). A system for example-based programming. In *Proceedings of CHI'89, Conference on Human Factors in Computing Systems*. April 30 - May 4, 1989. Austin, Texas.