NISTIR 5720

User's Guide to

"SuperFit"

Modeling Software for CMM Probe Lobing

Theodore H. Hopp

U.S. DEPARTMENT OF COMMERCE
Technology Administration
National Institute of Standards and Technology
Manufacturing Engineering Laboratory
Manufacturing Systems Integration Division
Gaithersburg, MD 20899-0001

October 1995

U.S. DEPARTMENT OF COMMERCE
Ronald H. Brown, Secretary

TECHNOLOGY ADMINISTRATION
Mary L. Good, Under Secretary for Technology

NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY
Arati Prabhakar, Director

DISCLAIMER

The identification of commercial products is given only for the sake of completely describing the operation of SuperFit. In no instance does such identification imply recommendation by the National Institute of Standards and Technology, nor does it imply that the particular product identified is necessarily the best available for the described purpose.

User's Guide to "SuperFit" Modeling Software for CMM Probe Lobing

Theodore H. Hopp
National Institute of Standards and Technology
Gaithersburg, MD 20899

Introduction

"SuperFit" is a software package developed by the National Institute of Standards and Technology (NIST) that fits a model of probe lobing to coordinate data obtained by probing a calibration ball. It runs under the Microsoft Windows™ 3.1 operating system. This document describes how to use the SuperFit software. Information about the probe lobing model itself, including when it is applicable, is being published in a separate report.[1]

This document is written as a step-by-step guide for using SuperFit Version 0.3. It was written assuming you have a general knowledge of how to use the Windows system. (If you do not, see your Windows documentation for how to run the Windows Tutorial from the Program Manager help.) Several formatting conventions are used in this guide. Text that you type or that appears on the screen is printed in a sans-serif font, like this: FILE NAME. Menu commands are written with a vertical bar separating the main menu selection from the selection in its associated pop-up menu. Thus, FILE|OPEN FILE... means the OPEN FILE... selection in the pop-up menu that appears under the main menu item FILE. (The underscores indicate the keyboard combinations to which the menu items are keyed. For instance, FILE|OPEN FILE... can be activated by holding down the ALT key and typing an F, and then typing an O.) In the left margin of this guide are occasional graphical icons. These reproduce command icons appearing in SuperFit's main screen, as explained below. The icon above is the appearance of the SuperFit application when it is minimized on your Windows desktop.

Start SuperFit

If you have not already done so, install the SuperFit software on your system. See the Appendix for installation instructions. Start SuperFit as you would any Windows application. When it starts, the main window will appear, and will look something like Figure 1. If necessary, resize the window border so that all parts of the window are visible, as shown in Figure 1.

[1]Estler, Phillips, Borchardt, Hopp, Witzgall, Levenson, Eberhardt, McClain, Shen, and Zhang, "Error Compensation for CMM Touch Trigger Probes," *Precision Engineering*, to appear.

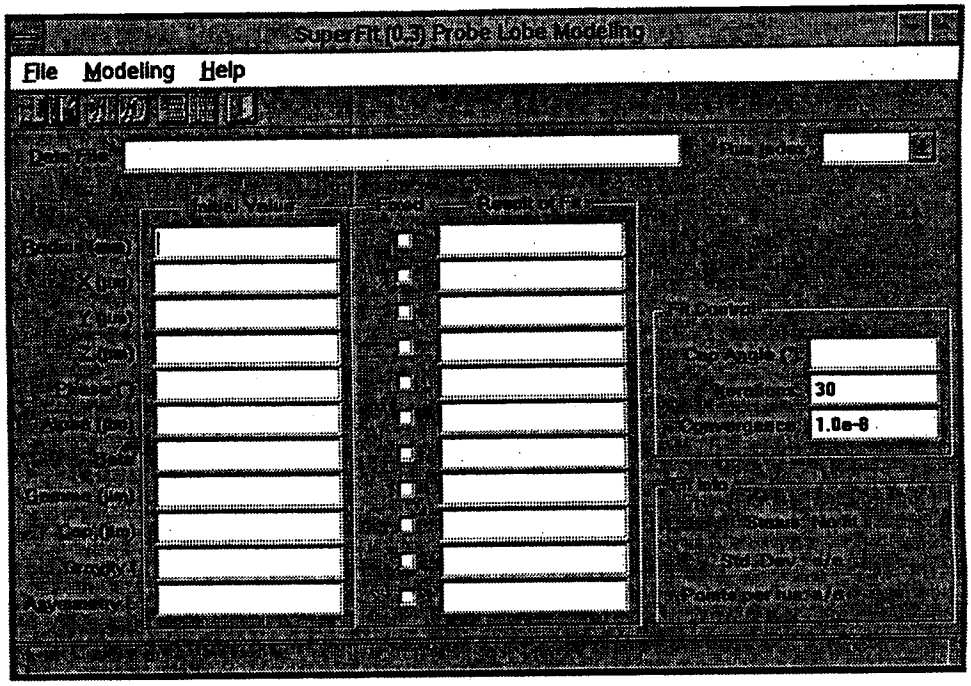

Figure 3 The main window for SuperFit.

Window layout

Across the top of the window is a menu bar with three items: FILE, MODELING, and HELP. Each of these is a pull-down menu containing commands that SuperFit can carry out. The FILE pull-down menu contains five commands:

FILE | OPEN FILE...
FILE | STORE RUNS...
FILE | WRITE REPORT...
FILE | EXPORT FITS...
FILE | EXIT

The MODELING pull-down menu contains two commands:

MODELING | FIT THIS RUN
MODELING | FIT ALL RUNS

The HELP pull-down menu contains two commands:

HELP | CONTENTS
HELP | ABOUT SUPERFIT

Only FILE | OPEN FILE..., FILE | EXIT and the HELP menu commands are enabled when SuperFit starts. The remaining commands are dimmed, indicating they are disabled.

Commands can be activated in several ways. You can select a command through the menu, as with any Windows program, using the mouse or ALT-*LETTER* combinations. Also, next to each command is a "hot key" keyboard combination you can use to activate the command directly. Finally, under the menu bar is a tool bar showing graphical icons. These icons represent the same commands that are available through the menu. You can click on an icon to carry out the corresponding command.

As you scroll through the menu commands or move the cursor over the icons, the status line at the bottom of the main window provides a short message about what the command associated with that menu item or icon will do if activated.

The main part of the window displays information fields about the model, most of which are blank when SuperFit starts. The information includes the name of the file containing the data being fit, the run index within the file (a file may contain several sets of data), initial values and fit results for model parameters, fit control data, and information about the fit.

As you proceed, you will see how the various parts of the window are used, and how commands become enabled. The first task is to open a file.

Open a File

Before proceeding further, you must have stored on disk a text file containing data from a probe measuring a calibration ball. The format expected by SuperFit is described in the Appendix. Open a file by selecting FILE|OPEN FILE... from the menu, by pressing CTRL+O on the keyboard, or by clicking on the corresponding icon. As with most Windows programs, the trailing dots indicate that a "dialog box" will appear asking you for more information. In this case, the dialog box will be a form asking you about the data file you want the program to use. The form is shown in Figure 2. The form has three main areas for you to enter information:

FILE	General information about the file to be opened
DIMENSIONS	Physical configuration of the probe used to acquire the data in the file
STYLUS ORIENTATION	Relationship of the stylus to gravity.

We will describe each of these in turn.

File information

The fields grouped in the FILE box provide general information about the data file to be opened. There are three fields:

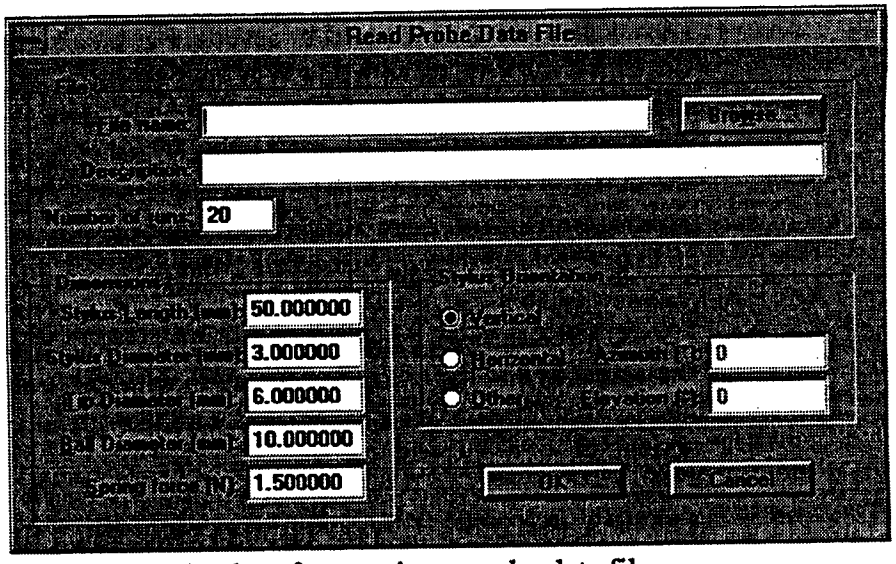

Figure 2 Dialog box for opening a probe data file.

FILE NAME—To identify the file, you enter the file name. If you press the button labeled BROWSE, a standard Windows file browser box will pop up. Select a file and the name will be entered into the FILE NAME field.

DESCRIPTION—You can also enter a short description (up to 80 characters) for the data in the file. This description will appear in any reports you generate later.

NUMBER OF RUNS—set this field to the number of data sets in the file. (If this is set too low, SuperFit will ignore the additional data sets; if it is set too high, SuperFit will report that it cannot read the file.)

Dimensions

The fields grouped in the DIMENSIONS box describe the physical configuration of the probe and calibration ball. These parameters are used to compute starting values for the model parameters to start a fit. SuperFit suggests values that correspond to a common configuration used at NIST. You are likely to be using a different configuration. All the fields are self-explanatory except for SPRING FORCE. This is a probe adjustment for which you will probably not have direct information. The suggested value of 1.5 N corresponds to about 6 grams-force to trigger a typical probe with a 50 mm stylus.

Orientation

The third group of fields, STYLUS ORIENTATION, relates the stylus to gravity. You should check that the correct choice (VERTICAL, HORIZONTAL, or OTHER) is selected. If you select HORIZONTAL, enter the azimuth of the stylus. The azimuth is the polar angle of the stylus in the X-Y plane. (SuperFit assumes that gravity is perpendicular to the X-Y plane and that the Z axis points up.)

If you select OTHER, enter both the azimuth and the elevation of the stylus from a straight-down vertical orientation.

Although the data grouped under DIMENSIONS need not be accurate (since they are only used to establish an initial guess), the orientation data must be correct. The probe lobing model is defined in terms of a stylus-and-probe coordinate system, and SuperFit uses the orientation data to transform the point coordinates into the correct system as it reads the file. With the wrong orientation, SuperFit will not be able to compute an accurate model.

When you are done entering information, press the OK button. SuperFit will start reading the data. As this can be a long process for large data sets, a progress gauge, such as that shown in Figure 3, will display. If you press CANCEL, reading of the data file will be stopped and any data read to that point will be discarded.

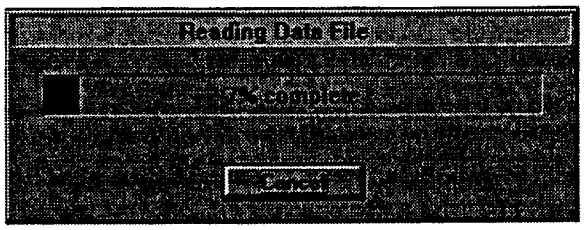

Figure 5 Progress indicator box.

When the data have been read, SuperFit returns to the main screen. At this point, the STORE RUNS command and the modeling commands are enabled, since they are applicable as soon as a file is open.

Save Runs

If you load a file that has several runs of data, you may wish to extract these into separate files. This is convenient, for instance, if you wish to then load the individual runs into a spreadsheet for analysis.

You can do this with the FILE | STORE RUNS... command (CTRL+S), which is available any time when a data file is loaded. When you extract data, SuperFit will ask you to identify a file name to store the data. SuperFit will add trailing indices to the file name of the form "_*n*" where "*n*" is the run number. For instance, if you enter a file name of "PROBE.DAT" and you have loaded a data file with ten runs, SuperFit will generate files named "PROBE_1.DAT" through "PROBE_10.DAT."

Only the coordinate data are written to the files. The data are in text format, one point per line, with values delimited by tabs.

Fit the Model

Once a data file has been read in, you can fit a model to the data. We first look at fitting a single run (that is, one measurement of a calibration sphere). We then explain how you can fit all the runs with one command. Before you fit, you should examine the values in two groups of fields in the main window, INITIAL VALUE and FIT CONTROL. You should also decide which model parameters should be varied by SuperFit and which should be held constant during the fit.

Initial values

You should check that the initial estimates set by SuperFit are reasonable. These values may affect the ability of the fitting routine to converge properly. You can modify any of the values at your discretion. (With experience you may be able to set initial values better than can SuperFit.) In particular, SuperFit assumes that the points are taken in a coordinate system roughly centered on the calibration sphere. If the sphere is not at the origin, you should set the initial values appropriately.

All of the initial values except for the phase are calculated by SuperFit from the dimension and orientation data you supplied when you opened the file. However, those data had no information about the phase of the lobing pattern with respect to machine coordinates. Therefore, SuperFit sets the phase angle to zero, an arbitrary value. If you have reason to believe another angle may be more appropriate (such as a prior model of the same probe), setting the phase to that value will greatly help convergence.

To the right of the initial guess fields is a column of boxes labeled FIXED at the top. If an "x" appears in a box, the corresponding fit parameter will be held fixed by SuperFit. (That is, it will not be varied as SuperFit computes the lobing model.) By default, all parameters are varied, except for the gravity effects parameter when the probe is vertical. If you wish to lock any model parameters to your initial guess, click on the corresponding box. (You can also use the TAB key on the keyboard to move control to a box and press the space bar to toggle the "x".)

Fit control

You should also check the three FIT CONTROL parameters on the right side of the main window: CAP ANGLE, ITERATIONS, and CONVERGENCE. The CAP ANGLE field should be set to the half-angle of the friction cone between the probe tip and the calibration sphere. The value of 11° set by SuperFit is reasonable for a ruby tip and a steel ball. The ITERATIONS field will limit the iterations that SuperFit will use in trying to reach convergence. The suggested value of 30 is adequate for a good initial guess, but you may need to set it to larger values (perhaps several hundred) if the data do not fit the model well. The CONVERGENCE field is the accuracy of the computed solutions. It should not be larger than 10^{-6} unless you are having serious problems getting the model to converge.

Fitting one run

Once you have reviewed the initial guess and set the fit control parameters, you can fit the run by issuing the MODELING|FIT THIS RUN command (CTRL+F) from the menu or clicking on the corresponding icon. SuperFit will display a progress gauge similar to that of Figure 3. It uses the iteration limit to estimate how much work is left, so the fit may converge well before progress reaches 100%.

If you press CANCEL, fitting stops and the fit results are discarded for that run. It may take SuperFit a few seconds to respond to a cancel signal.

Once at least one fit has been computed for a data file, the FILE|WRITE REPORT... and FILE|EXPORT FITS... commands are enabled. These commands are discussed below.

Fit information

After the fit is computed, you can look at the status information grouped in the FIT INFO box at the lower right of the main window. It will show the result code from fitting (that is, why SuperFit stopped working), as well as the standard deviation of the residual lobing and the number of points.

If the STATUS field displays "CONVERGED," the fitting routine converged successfully. However, it may have converged to a local minimum. Indications of false convergence are negative values for alpha, beta, gamma, or cap. If you have these conditions, change the starting guess for the phase by 60° and recompute the fit. If the STATUS field displays "ITERATION LIMIT," increase the iteration limit field and try the fit again. If you canceled the fitting, the STATUS field will display "USER BREAK." Any other status is the result of an internal error in SuperFit and should be reported to NIST.

The standard deviation is computed using an unbiased estimate: the sum of squared residuals is divided by the number of points minus the number of non-fixed model parameters before taking the square root.

If the data file has data for more than one run, SuperFit allows you to fit each run individually using the method described above. Click on the arrow next to the RUN INDEX field to display a list of all the runs in the data file, from which you can select the run you wish to use. You can also use the up and down arrow keys on the keyboard to go through the list of runs.

Fitting all runs

As an alternative to working with one run at a time, SuperFit allows you to fit all runs at once. After you have set the initial guess and fit control fields, you can fit all the runs by selecting MODELING|FIT ALL RUNS (CTRL+A) from the menu or clicking on the corresponding icon. This command will use the same initial guess for all runs. If you press "CANCEL" during fitting, fits

already computed are kept; only the fit for the run currently being computed will be discarded.

Once all the fits are computed, you can review the status of each fit by going through the run index list. If any run stopped at the iteration limit, or appears to have reached false convergence, you can use the FIT ONE RUN procedure above to recompute the fit for that one run alone.

Write a Report

Once you have fit the data, you can write a formatted report of the results. Select FILE|WRITE REPORT... (CTRL+W) from the menu or click on the corresponding icon. SuperFit will display a standard file selection dialog box. Enter a file name and press OK, and a report will be written to that file.

Export Fits

Similarly, you can write fit results into a file suitable for import into a spreadsheet. You do this by selecting FILE|EXPORT FITS... (CTRL+P) from the menu or by clicking on the corresponding icon. Enter a file name in the file selection dialog box and press OK. The file will contain one header row of parameter labels and one row of parameter values for each fit you have computed. This is a text file with comma-delimited fields and with labels enclosed in double quotes (").

Exit SuperFit

When you are finished, exit SuperFit by selecting FILE|EXIT (ALT+F4) from the menu or clicking on the corresponding icon. If you have calculated fits but have not written a report or exported the fits to a spreadsheet file, SuperFit will ask you to confirm that you want to exit at that point. If you do, press OK to go ahead; otherwise, press CANCEL to return to SuperFit.

If you finish working with one data file but want to work with another, you do not need to exit SuperFit. Just select the FILE|OPEN FILE... command. The current file will be unloaded and a new one read in. As with exiting, SuperFit will ask you if you wish to discard existing fits if you have not written a report or not exported the fit results.

Getting Help

The HELP menu contains two commands that you can use to find out more about SuperFit. The HELP|CONTENTS (ALT+F1) command provides on-line documentation similar to this manual. The HELP|ABOUT SUPERFIT command displays a dialog box showing information about the version of SuperFit you are running. This information is useful if you need to report errors to NIST.

APPENDIX—INSTALLATION AND SUPERFIT FILES

Installation

SuperFit is distributed as a self-extracting archive on a floppy disk. To install SuperFit, first create a directory on your hard drive and copy the file INSTALL.EXE from the floppy to that directory. Then run INSTALL.EXE. (Do not run INSTALL.EXE directly from the floppy; it will attempt to decompress the files onto the floppy and run out of room.) INSTALL.EXE will create five files in the directory. The SuperFit program file is named SUPERFIT.EXE. MANUAL.WP is this User's Guide, in WordPerfect 6.1 format. The three other files, BC450RTL.DLL, BIDS45.DLL, and OWL250.DLL are dynamic link libraries used by SuperFit.[2] The three DLLs are used by many Windows programs, and may already be present on your system (usually in the Windows directory). If you already have them elsewhere on your system, you may wish to delete them from the SuperFit directory to save disk space. You may also delete INSTALL.EXE.

Program Item

You can create a program group and program item for SuperFit in Windows Program Manager. See your Windows documentation for instructions on how to do this. This allows you to start SuperFit by clicking on an icon. Otherwise, you will have to select FILE | RUN from the Program Manager (or File Manager) in Windows and enter the path to SUPERFIT.EXE.

Data file format

SuperFit will only read coordinate data formatted in the way used at NIST for our experiments. Other data file formats will be supported in future versions of SuperFit. The data must be in a text file with data for one point per line. Each line must contain seven numbers, in the following order:

1. Run number. Each file can contain data from many repetitions of a probe measurement. However, all data must be for the same configuration and stylus orientation. Also, all runs must have the same number of points. Run numbers are integers starting at one and increasing by one.
2. Point number within the run. Point numbers are integers starting at one and running consecutively to the number of points per run.
3. Point declination angle. This number is ignored by SuperFit, but must be present.
4. Point azimuth angle. This number is ignored by SuperFit, but must be present.
5. The *X* coordinate of the point.
6. The *Y* coordinate of the point.
7. The *Z* coordinate of the point.

[2]The three .DLL files are proprietary property of Borland International, and are distributed by NIST with permission.