

Figure 8: Thirty-Five Pins of a 68 pin Cerquad Package

### 1.4.5 Right Angle Connector

The example file `30pin.inp` from Teradyne Connection Systems Inc. shown in Figure 9 is a PCB right angle connector plugged into its thirty corresponding pins. The V-shaped clamping portion and right angle bend make up the connector as shown in the single pin illustration of Figure 10. In reality, the pin protrudes through the V-shaped clamp, but that section conducts no current and is not included.

Each of the thirty pins is a different shade of gray. The white portions of the V are also part of each pin but appear white due to the limits of the shading algorithm.

### 1.4.6 Photodetector

The example file `msm.inp` shown in Figure 11 is a metal-semiconductor-metal photodetector from NASA Langley Research Center and University of Virginia. These devices consist of a set of closely spaced, interdigitated electrodes deposited on an optically active semiconductor layer. Every other electrode is biased to the same potential, so that any two adjacent electrodes are at opposite electrostatic potentials. Because the electrodes are closely spaced, the electric fields between them are quite strong with application of low to moderate bias voltages. As a result, any carriers generated by absorption of light between the fingers are transported to the electrodes extremely rapidly giving fast response to short pulses of light.

## 2 Running FastHenry

The basic form of the FastHenry program command line is

```
fasthenry [<input file>] [<Options>]
```

Usually only the `input file` is specified. For example, the command

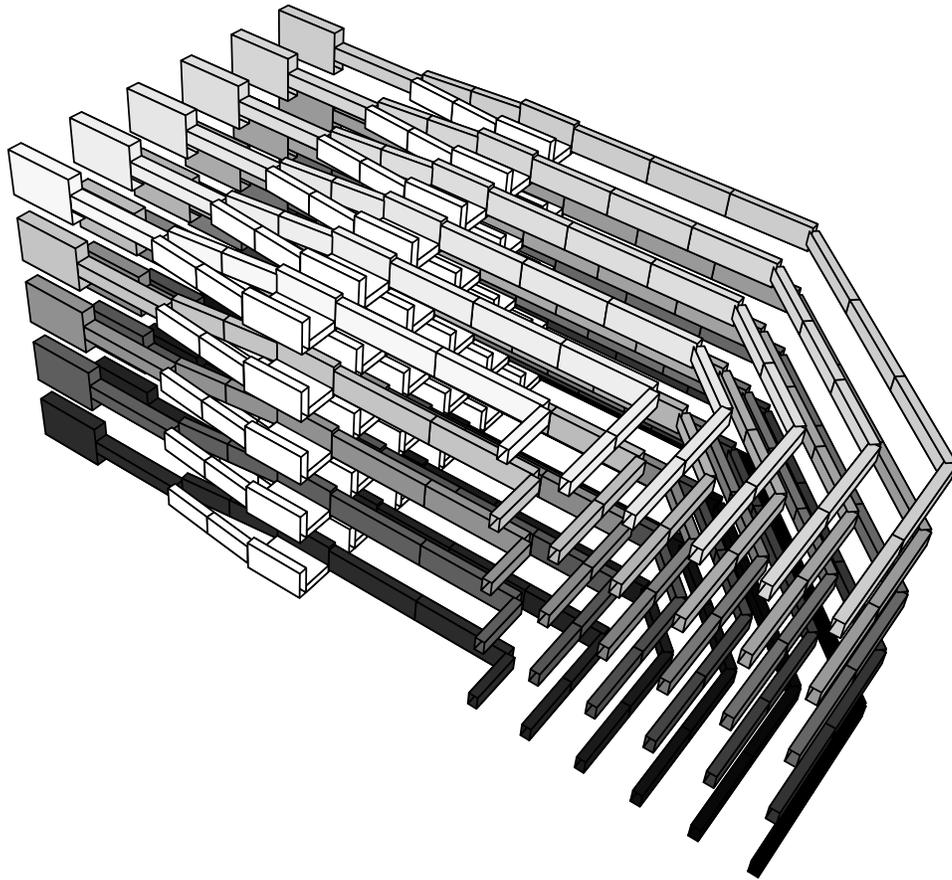


Figure 9: Thirty pin right angle connector

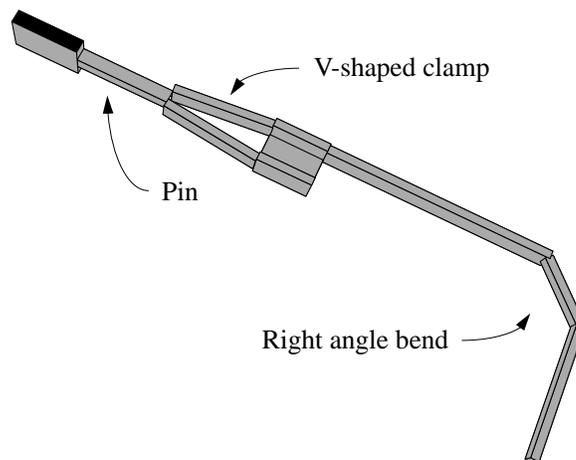


Figure 10: One pin of the thirty pin right angle connector

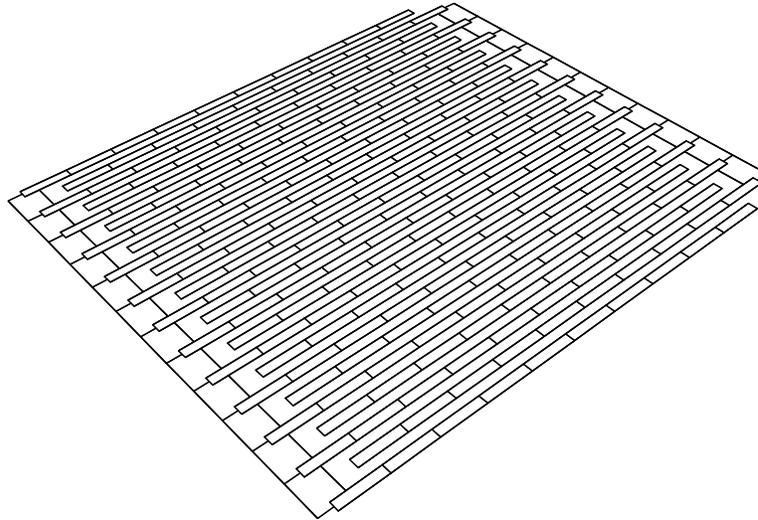


Figure 11: Metal-semiconductor-metal Photodetector

```
fasthenry pin-connect.inp
```

runs `fasthenry` on the example pin connect structure.

Information about the input file and other FastHenry information are sent to the standard output. The impedance matrices for the frequencies specified in the input file will be placed in the text file `Zc.mat`. This file also lists the correspondence between columns in the impedance matrix and the ports specified in the input file. The source file `ReadOutput.c` described in Section 2.2 is provided as a sample program for reading the output file for postprocessing.

## 2.1 Example Run

This section contains a sample run of FastHenry for the 30 pin connector example, `30pin.inp` shown in Figure 9. Comments describing the output appear along the right margin in addition to the FastHenry command line option which would change the described setting.

```
prompt % fasthenry 30pin
Running FastHenry 2.0 (10Aug94)
  Date: Wed Aug 10 16:31:35 1994
  Host: lyapunov
Solution technique: ITERATIVE          <-- use GMRES (-s)
Matrix vector product method: MULTIPOLE <-- and multipole algorithm (-m)
  Order of expansion: 2                 <-- order of multipole (-o)
Preconditioner: 0N                      <-- GMRES preconditioner (-p)
Error tolerance: 0.001                  <-- relative tol (-t, -b)
Reading from file: 30pin
Title:
**30 pin, right angle connector**      <-- first line from input file
```



```

    Multipole setup 11.94
    Scanning graph  0.01
    Form A M and Z  0.06
    form M'ZM       0
    Form precondition 1.36
    GMRES time      75.92
Total:              90.8
90.980u 0.640s 1:36.62 94.8% 321+2448k 0+0io 4pf+0w
prompt %

```

## 2.2 Processing the Output

### 2.2.1 The impedance matrix

The file `Zc.mat` is a text file containing the impedance matrices for the frequencies requested. The file `ReadOutput.c` in directory `src/misc/` is an example program for reading the text file for whatever processing is necessary. It contains the function `ReadZc()` which reads from a file and returns a linked list of impedance matrices and their corresponding frequencies. See the source file for more details. This function can be extracted and included in whatever program the user desires.

The function `main()` is provided as an example use of `ReadZc()`. For each of the matrices, it divides the imaginary part by the frequency to give the matrix  $R + jL$  and then dumps the result to the standard output.

`ReadOutput.c` can be compiled by typing

```
cc -o ReadOutput ReadOutput.c
```

Here is a sample of its output after processing `Zc.mat` produced by running `fasthenry` on example file `onebargp.inp`:

```

prompt % ReadOutput Zc.mat
Not part of any matrix: Row 2 : nodein to nodeout
Not part of any matrix: Row 1 : nb1 to nbout
Reading Frequency 10000
Reading Frequency 100000
Reading Frequency 1e+06
Reading Frequency 1e+07
Reading Frequency 1e+08
freq = 1e+08
Row 0: 0.00112838+3.50478e-08j -2.21062e-05-7.0003e-09j
Row 1: -2.23118e-05-6.99564e-09j 4.83217e-05+2.02958e-08j
freq = 1e+07
Row 0: 0.00112837+3.50478e-08j -2.21055e-05-7.0003e-09j
Row 1: -2.2311e-05-6.99564e-09j 4.83217e-05+2.02958e-08j
freq = 1e+06
Row 0: 0.0011272+3.50501e-08j -2.20335e-05-7.00045e-09j
Row 1: -2.22379e-05-6.99578e-09j 4.83168e-05+2.02958e-08j
freq = 100000

```

```

Row 0: 0.00106093+3.51782e-08j -1.7938e-05-7.00794e-09j
Row 1: -1.80765e-05-7.00341e-09j 4.80425e-05+2.02964e-08j
freq = 10000
Row 0: 0.000955377+3.58555e-08j -1.25489e-05-7.01913e-09j
Row 1: -1.25828e-05-7.01518e-09j 4.75474e-05+2.03045e-08j
prompt %

```

### 2.2.2 Creating Equivalent Circuits

Two approaches for generating spice equivalent circuits are available:

#### Approach 1: An equivalent circuit for a single frequency.

- Run `fasthenry` for EXACTLY one frequency.
- Compile (with some C compiler) the file `MakeLcircuit.c` in the directory `fasthenry-3.0/src/misc/`

```
cc -o MakeLcircuit MakeLcircuit.c -lm
```

- Run `MakeLcircuit` on the `Zc.mat` produced by `FastHenry`

```
MakeLcircuit Zc.mat > my_circuit.spice
```

This produces a circuit where nodes 0 and 1 correspond to the plus and minus nodes of the first `FastHenry` port (also called `.external`). 2 and 3 correspond to the second port. 4 and 5 the third, etc.

#### Approach 2: A circuit which models the frequency dependent resistances and inductances

This approach generates a reduced order model for the system using the arguments “-r n -M” to `FastHenry`. This gives an equivalent circuit which is valid for a range of frequencies.

- Run `FastHenry`: `fasthenry -r n -M myinput.inp`  
Where “n” is replaced with the desired order, say 20. This produces the file `equiv_circuitROM.spice` containing a single spice subcircuit.
- Include the subcircuit in a spice input file and connect devices to it.

The subcircuit will be named `ROMequiv` and its port nodes are `p0` and `m0` for the plus and minus terminals of port 0, `p1` and `m1` for port 1, etc. See the header of the `equiv_circuitROM.spice` file for more description. Any suffix specified with “-S” will be appended to both the subcircuit name, `ROMequiv`, and the subcircuit filename, `equiv_circuitROM.spice`.

Comments: Approach 1 will not model frequency dependent resistance and inductance since it gives an R and L at the single specified frequency. Approach 2 will model the full

frequency dependent effects up to some frequency where that frequency is higher for a higher chosen order,  $n$ . The reduced order model of Approach 2 is converted to a circuit from its state-space form through capacitors, resistors, and VCCSs so insight can only be gained by simulating the subcircuit in spice rather than looking inside the subcircuit file. An example is provided below. For a description of reduced order modeling for inductance computation see the paper `tchmt-epep94.ps` in the `pub/fasthenry` directory at the `rle-vlsi.mit.edu` ftp site:

L. Miguel Silveira and Mattan Kamon and Jacob K. White, “Efficient Reduced-Order Modeling of Frequency-Dependent Coupling Inductances associated with 3-D Interconnect Structures”, *IEEE Transactions on Components, Hybrids, and Manufacturing Technology, Part B: Advanced Packaging*, Vol. 19, No. 2, May 1996, pp. 283-288.

### An example of Approach 2

The file `examples/pin-con7.inp` is 7 pins of a larger package shown in Figure 8. Each conductor is discretized into filaments to capture skin and proximity effects up to around  $10^{10}$  Hz. One can generate an equivalent circuit via Approach 2 with

```
fasthenry -r 20 -M pin-con7.inp -S_pin_con7_r20
```

which produces the file `equiv_circuitROM_pin_con7_r20.spice` which is an equivalent circuit for a 140th ( $20 * 7$ ) order model. To test its accuracy, consider computing impedance matrices at many frequencies with

```
fasthenry pin-con7.inp -S_pin_con7
```

which produces `Zc_pin_con7_r20.mat` which is the impedance matrix for frequencies  $1, 10, 10^2, \dots, 10^{12}$ . A spice file is included in the `examples` directory to test this subcircuit and put the results in `rom_check_con7_r20.out`:

```
spice3 -b -o rom_check_con7_r20.out rom_check_con7_r20.ckt
```

The spice file in `rom_check_con7_r20.ckt` has a spice `.include` line to include the subcircuit of file `equiv_circuitROM_pin_con7_r20.spice`. It then does an AC analysis of the response of the circuit when 1 Amp is applied to port 0 and 0 Amp to the other ports. The results correspond to column 1 of the impedance matrix,  $Zc$ .

A comparison of the accuracy of entry (1,1) of the impedance matrix is given in Figure 12 for both “-r 7” and “-r 20”.

As can be seen in the figures, in the order 7 reduced-order model the resistance matches the directly computed values (`Zc.mat`) up to about  $10^8$  Hz. The order 20 reduced-order model is more accurate, and represents the frequency-dependent resistance up to  $10^9$  Hz. The inductance in the example chosen does not have a strong frequency dependence.

Note that if Approach 1 were used, a single frequency would have to be chosen to generate an equivalent circuit. This would be less accurate but would probably reduce subsequent Spice simulation time.

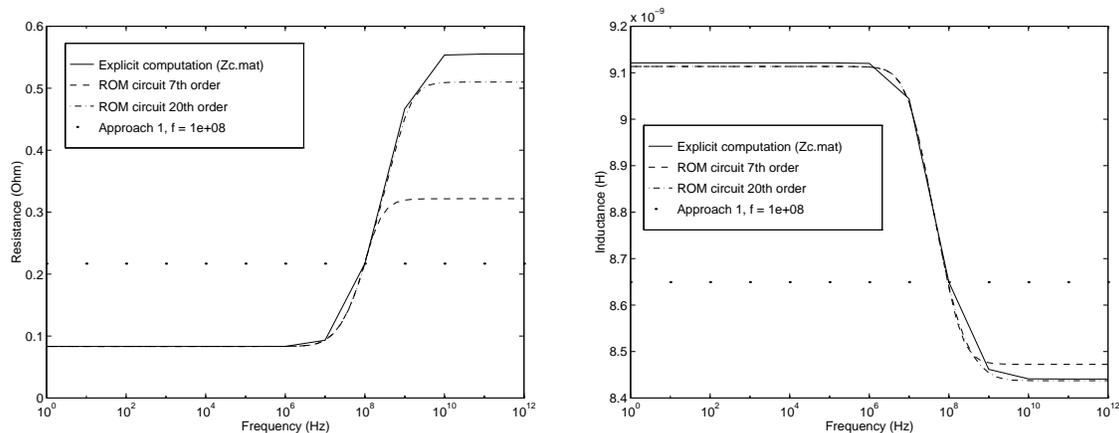


Figure 12: Comparison of impedance computation done via explicit solution at specified frequency points and reduced-order models

## 2.3 Command Line Options

This section describes using the command line options to change the defaults settings. All arguments are case insensitive.

- (dash) Forces input to be read from the standard input.

- s {`ludcomp` | `iterative`} - Specifies the matrix solution method used to solve the linear system arising from the discretization. `iterative` uses the GMRES iterative algorithm and `ludcomp` uses LU decomposition with back substitution. In general, GMRES is faster, however some speed up may be obtained using LU decomposition for problems with fewer than 1000 filaments. `iterative` is the default.

- m {`direct` | `multi`} - Specifies the method to use to perform the matrix-vector product for the iterative algorithm. `direct` forms the full matrix and performs the product directly. `multi` uses the multipole algorithm to approximate the matrix-vector product. For larger problems, the multipole algorithm can save both computation time and memory. `multi` is the default.

- p {`on` | `off` | `loc` | `posdef` | `cube` | `seg` | `diag` | `shells`} - Specifies the method to precondition the matrix to accelerate iteration convergence. Two classes of preconditioners are implemented in FastHenry. Local inversion preconditioners (`loc` and `posdef`) were used in prior releases and have been replaced by the sparsified-L preconditioners. Three types of sparsified-L preconditioners are available, (`cube`, `seg`, and `diag`). `cube` is the default and produces the best results, however under certain conditions can consume significantly more memory than `diag`. If there are multiple filaments in each segment, than `seg` is a middle ground between `diag` and `cube`, otherwise it is identical to `diag`. `shells` uses a current shell idea described in

Mattan Kamon, Byron Krauter, Joel Phillips, Lawrence T. Pileggi, Jacob White, "Two Optimizations to Accelerated Method-of-Moments Algorithms

for Signal Integrity Analysis of Complicated 3-D Packages”, *Proceedings of the IEEE 4th Topical Meeting on Electrical Performance of Electronic Packaging*, Portland, OR, October 1995, pp. 213-216.

The method is relatively new, and its performance is sensitive to the shell radius set by the `-R` option described below.

`-o n` - Specifies  $n$  as the order of multipole expansions. Default is 2. Choosing values less than 2 result in faster execution at the expense of poorer accuracy. Changes in the order of expansions should accompany changes in iteration error tolerance with the `-t` option.

`-l {n | auto}` - Specifies  $n$  as the number of partitioning levels for the multipole algorithm. `auto` chooses the level automatically and is the default.

`-f {off | simple | refined | both | hierarchy }` - Switches FastHenry to visualization mode only and specifies the type of FastCap generic file to make (for visualization ONLY!). `off` will produce no file and is the default. `simple` will produce a file named `zbuffile` from the segments defined in the input file. `refined` produces a similar file, named `zbuffile2`, but uses the segments after either user refinement specified with `-i` or required refinement necessary for accuracy of the multipole algorithm (`-l` option). `both` produces both files. One FastCap “panel” is created for each of the four sides along the length of the each segment. Reference planes are handled differently as described under the `-g` option below. See Section 3 for details on producing postscript. FastHenry will exit after creating the visualization files. `hierarchy` will dump a 2D representation of the nonuniform reference plane discretization hierarchy to the FIXED file `hier.qui` which can also be processed like the `zbuffile`.

`-g {on | off | thin | thick}` - controls appearance of the ground plane when using the `-f` option. `off` is the default and only four panels are produced for each plane, one for each of the edges. Thus only the outline of each plane is drawn. This makes generation of the postscript image much faster and also makes the planes transparent. No holes are visible, however. `on` or `thin` will draw all the overlapping segments of the ground plane as if infinitely thin. Note that using this option for a finely discretized plane may take a long time to generate a postscript image with `zbuf`. `thick` will completely draw the segments of the plane and will be even slower for generating postscript images.

`-a {on | off}` - `on` allows the multipole algorithm to automatically refine the structure as is necessary to maintain accuracy in the approximation. The structure will be refined whether or not the multipole algorithm is used. `off` prevents refinement and will produce a warning if the multipole algorithm is used and prevented from necessary refinement. `on` is the default and is equivalent to `-l auto`. Note that this is not refinement to reduce discretization error. See the `-i` option.

**-i n** - Specifies  $n$  as the level for initial refinement. This option allows the user to refine the structure if the input file is too coarse. It will divide each segment of the geometry into multiple segments so that no segment has a length greater than  $\frac{1}{2^n}$  times the length of the smallest cube which contains the whole structure. The default is 0 (no refinement). This option is rarely used since the multipole algorithm will refine as it needs and discretization errors of this sort are usually small.

**-d {on | off | mrl | mzmt | grids | meshes | pre | a | m | rl | ls}**  
 - dump certain internal matrices to files. The format of some of the files can be specified with the **-k** option. **on** dumps the  $M, R, L, MZM^t$  and preconditioner matrices. **off** dumps none and is the default. **mrl** dumps the M, R, and L matrices. **mzmt** dumps the  $MZM^t$  matrix for  $w = 1$ . **grids** dumps matrices for viewing the current distribution inside each reference plane (only in matlab format). **meshes** is a matlab file for viewing the KVL meshes chosen by FastHenry. **pre** dumps the preconditioner in a sparse matrix text format. **ls** dumps the sparsified  $L$  matrix in sparse text format. **a** is the branch incidence matrix,  $A$ . **m** can be used to dump only the mesh matrix  $M$ . Similarly, only R and L can be dumped with **rl**. The right-hand-side vector,  $V_s$ , is dumped to the file **b.mat** whenever anything but **off** is specified.

**-k {matlab | text | both}** - Specifies type of file to dump with the **-d** option. **matlab** dumps the files as M.mat, MRL.mat, MZMt.mat, and A.mat in a format readable by matlab. **text** saves the files M.dat, L.dat, R.dat, MZMt.dat, and A.dat as text. **both** saves files in both formats.

**-t rtol, -b atol** - Specifies the tolerance for iteration error. FastHenry calculates each column of the impedance matrix separately. The iterative algorithm will stop iterating when both the real and imaginary part of each element,  $x_k$ , of the current column being calculated satisfy

$$|x_k^{i-1} - x_k^i| < rtol * (|x_k^i| + atol * (\max_j |x_j^i|)) \quad (1)$$

where  $i$  is the iteration number. The defaults are  $rtol = 10^{-3}$  and  $atol = 10^{-2}$ . In essence, this causes GMRES iterations to stop if every element of solution vector has changed by less than the fraction  $rtol$  since the last iteration. If, however, element  $x_k$  is smaller than  $atol$  times the largest element in the solution vector, then the stopping criteria is less stringent. Note that real and imaginary parts are treated separately.

**-c n** -  $n =$  maximum number of iterations to perform for each solve (column). Overrides the default of 200.

**-D {on | off}** - Controls the printing of debugging information. **off** is the default. **on** will cause FastHenry to print more detailed information about the automatic partitioning level selection, memory consumption, preconditioner calculation, and convergence of the iterates. The Matlab file **Ycond.mat** is also produced which contains the admittance matrices at each frequency and also the norm of the residual at each step of the GMRES algorithm.

**-x portname** - Specifies that only the column in the admittance matrix specified by **portname** should be computed. Multiple **-x** specifications can be used. In this case, the output file **Zc.mat** will contain the requested columns of the admittance matrix instead of the impedance matrix. The portname is specified in the input file with the **.External** keyword. The primary uses of this option are either to exploit symmetry or to compute single columns for observing current distribution. For instance, if a 10 pin package is symmetric about some axis, then only five columns need be computed. The user is then responsible for forming the  $10 \times 10$  impedance matrix from the  $10 \times 5$  result. Also, if the **-d grids** option is used, the port which is the source of current would be specified with this option.

**-S suffix** - This adds the string **suffix** to all filenames for this run. For instance, **-S \_blah** will produce the output file **Zc\_blah.mat**.

**-r order** - Specifies a reduced order model of the system as output. The size of the model will be **order**\*number-of-ports. A matlab file **rom.m** will contain the A,B,C,D matrices representing the state-space reduced order model. More useful is the file **equiv\_circuitROM.spice** which is a spice3 compatible subcircuit which can be included for circuit simulation.

**-M** - If **-r** is specified with a nonzero order, then this option will cause FastHenry to exit after generating a reduced order model.

**-R radius** - If **-p shells** is specified, then this specifies the radius of the shells to use.

**-v** - Regurgitate internal representation of the geometry to stdout in the input file format. Good for debugging. User can compare output to the original input file. Also, by altering the file **regurgitate.c**, the user can also call functions to translate and reflect the geometry for use as another input file.

## 2.4 Discretization Error Analysis

This section gives some guidelines and describes an experimental approach to determining an appropriate discretization to model skin and proximity effects in long, thin conductors defined with FastHenry segments. (For a description of discretization for reference planes, see “nonuniform\_manual.ps”.)

To begin, assume it is desired to compute the impedance matrix for the pin connect structure of Figure 8 for frequencies up to  $10^8$  Hz.

### 2.4.1 The DC case

If only the DC case were of desired, it could be computed rapidly since there is no skin effect. The number of filaments per segment for the entire structure is set to one by setting **nhinc=1** and **nwinc=1** with the line

```
.DEFAULT Z=85. H=8.5 W=24. nhinc=1 nwinc=1
```

In order to extract a nonzero value for the inductance, the frequency is set small, but nonzero

```
.freq fmin=1e0 fmax=1e0 ndec=1
```

Since there are only a few hundred filaments, there is no advantage to the multipole algorithm so LU decomposition can be used to save a few seconds:

```
fasthenry pin-connect.inp -sludecomp -S _DC
```

The output will be in file `Zc_DC.mat`

### 2.4.2 The highest frequency case

Next, consider deciding how many filaments are needed to accurately compute the resistance and inductance at the highest frequency point of interest,  $f = 10^8$  Hz. A good rule of thumb is to choose the discretization such that the width of the smallest filament is roughly equal to the skin depth. For this geometry,  $\rho = 0.0238$  ohm-mil or  $\sigma = 1.641 \times 10^6$  mho/m giving a skin depth of

$$\delta = \sqrt{\frac{1}{\pi f \mu \sigma}} = 1.55 \text{ mils.} \quad (2)$$

To observe the error involved with various discretizations, we will look at the impedance matrix for two typical adjacent segments from the `pin-connect.inp` example:

```
***Two segments from: PACKAGE INDUCTANCE - 35 of 68 pins***
*   For experimenting with skin/proximity effect modeling
*
. UNITS MILS
. DEFAULT RHO=.0238
. DEFAULT Z=85. H=8.5 W=24. nhinc=3 nwinc=5 rh=2 rw=2
*
* BEAM 16
*
N16B Y=100. X=387.5
N16C Y=100. X=493.
E16C N16B N16C
. EXTERNAL N16B N16C port16
*
* BEAM 17
*
N17B Y=50. X=387.5
N17C Y=50. X=493.
E17C N17B N17C
```

```
.EXTERNAL N17B N17C port17
*
.freq fmin=1e1 fmax=1e12 ndec=3

.END
```

The result will be a  $2 \times 2$  complex matrix for each frequency point,

$$\mathbf{Z}(\omega) = \begin{bmatrix} R_{11}(\omega) + j\omega L_{11}(\omega) & R_{12}(\omega) + j\omega L_{12}(\omega) \\ R_{21}(\omega) + j\omega L_{21}(\omega) & R_{22}(\omega) + j\omega L_{22}(\omega) \end{bmatrix} \quad (3)$$

If we choose  $nhinc = 3$  and  $nwinc = 5$  with a default ratio of adjacent filaments equal to 2 ( $rh = rw = 2$ ), then the width of the smallest filament on the right and left of the segment will be  $1/10$  the width  $((1 + 2 + 4 + 2 + 1)^{-1})$ , or 2.4 mils. The thickness of the filaments on the top and bottom of the segment will be  $1/4$  of the height  $((1 + 2 + 1)^{-1})$  or 2.1 mils.

To run fasthenry,

```
fasthenry pin-con2seg.inp -sludecomp -aoff -S _3x5_2
```

In this case, since there are so few filaments, LU decomposition is slightly faster. Also, automatic refinement is prevented with the `-aoff` option since the multipole algorithm is not being used and this problem is essentially two-dimensional.

The smallest filaments are slightly above the skin depth and thus the results are inaccurate at  $f = 10^8$  as shown by the dotted lines in Figure 13. The next logical step might be to increase the discretization to  $nhinc \times nwinc = 4 \times 6$  however using even numbers of filaments is not efficient since the center of the segment, where current is relatively constant, will be divided into 4 separate filaments. For the odd case, the center is modeled by only one filament. Instead, consider changing the ratio of adjacent segments to  $rh = 4$ ,  $rw = 3$ . Now, the width of the smallest filament on the left and right sides of the segment will be  $1/17$  of the total width  $((1 + 3 + 9 + 3 + 1)^{-1})$  or 1.41 mils and the thickness of the smallest filament on the top and bottom, 1.42 mils. This is much closer to the skin depth and gives better results as shown by the dash-dot lines in Figure 13.

Another possibility is to reduce the number of filaments below  $3 \times 5$  but further increase the ratio of adjacent filaments. This change would result in faster execution since there would be fewer filaments, however some accuracy at midrange frequencies would be sacrificed. For instance, consider maintaining the width of the smallest filaments by changing the discretization above to  $3 \times 3$  with  $rh = 4$ ,  $rw = 15$ . In this case there are too few filaments to model the decay of current density from the outside edge to center for  $f \in [10^6, 10^8]$  as shown by the dashed line in Figure 13.

After the user has decided which error for the above cases is tolerable for a given problem, the discretization can then be used for the full `pin-connect.inp` example. The default  $nhinc$  and  $nwinc$  must be set with

```
.DEFAULT Z=85. H=8.5 W=24. nhinc=3 nwinc=5 rh=4 rw=3
```

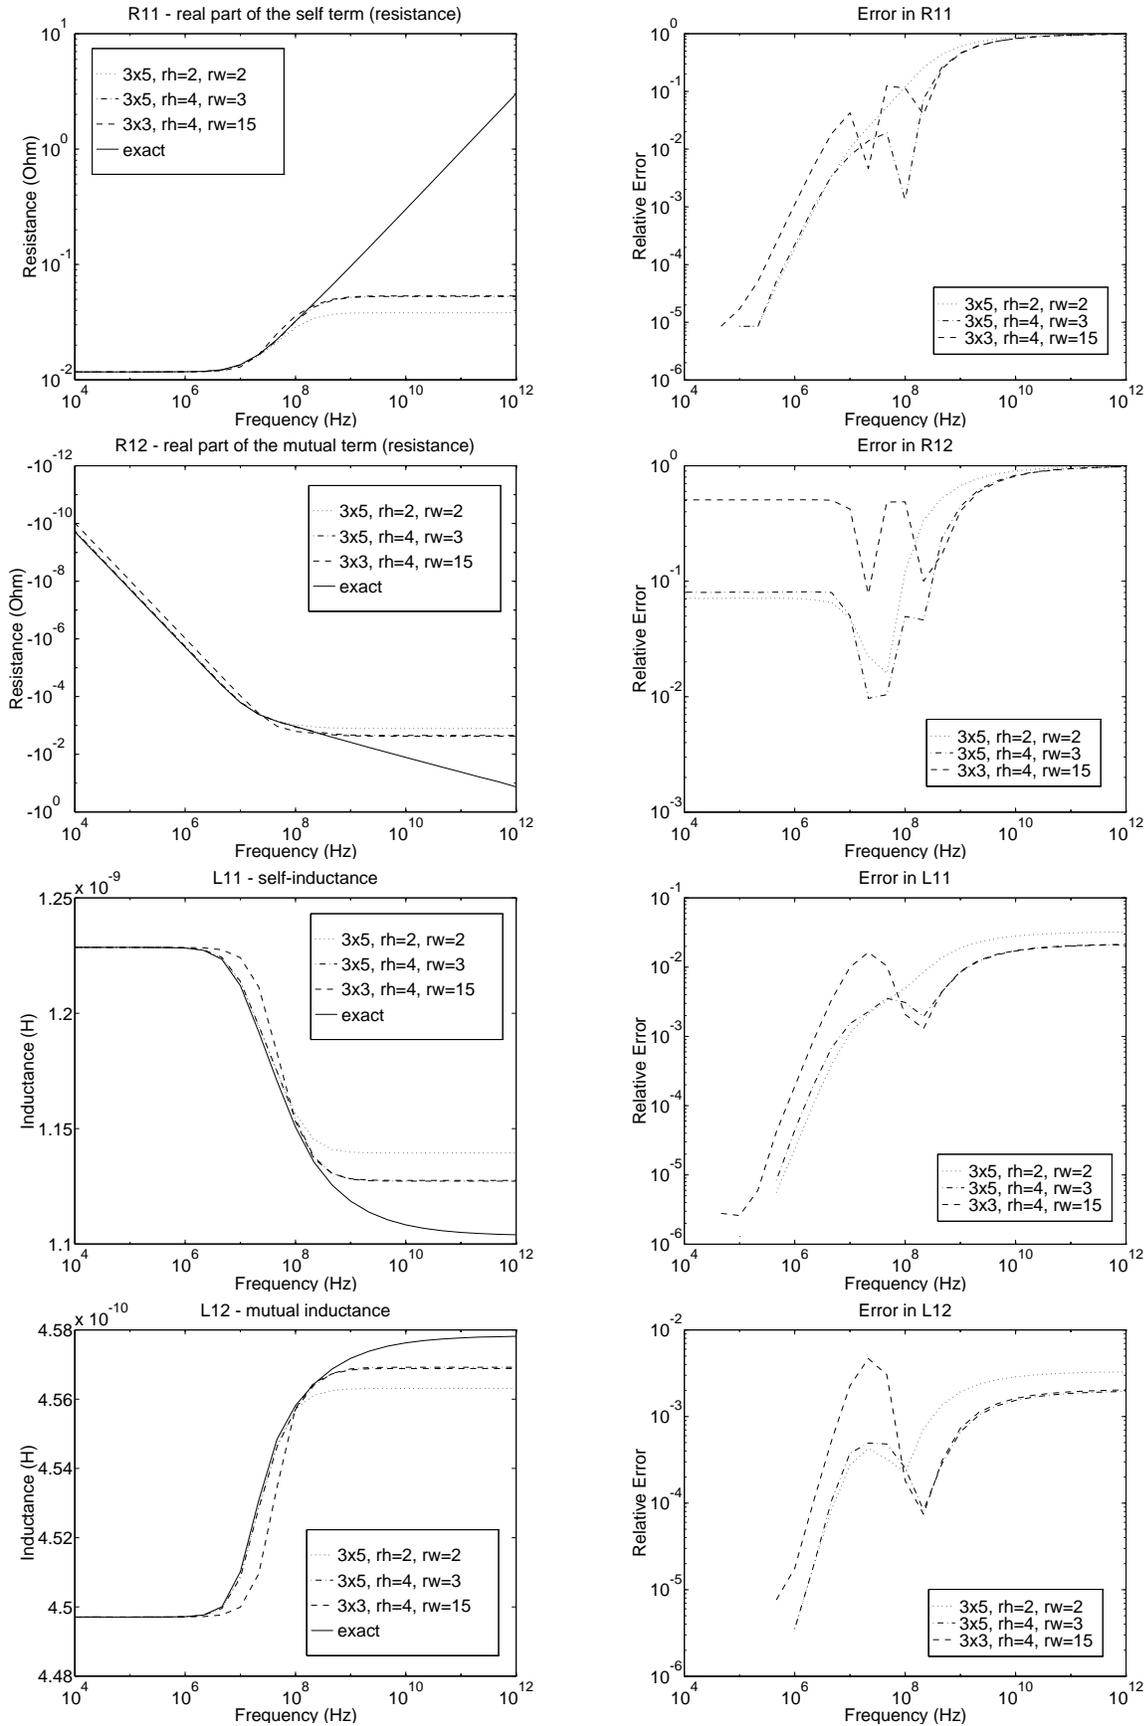


Figure 13: Values and errors for the impedance matrix as a function of frequency for the 2 segment experiment

and the frequencies of interest with

```
.freq fmin=1e0 fmax=1e12 ndec=3
```

The frequency setting above was used to generate Figure 13. Note that with the sparsified-L preconditioner, used by default, the low frequency impedance matrix is computed in very few iterations. Thus, while the low frequency case could have been computed with a coarser discretization as in Section 2.4.1, there is not much added execution time in computing it at a finer discretization.

The exact values in Figure 13 were generated using a  $21 \times 25$  discretization with  $rh = rw = 2$ . It was run with

```
fasthenry pin-con2seg.inp -aoff -mdirect
```

The multipole algorithm would require space be divided, and thus the filaments also be divided, tripling the number of filaments in this case. Since the original number of filaments is small and, again, this is essentially a two-dimensional problem, we can avoid both dividing space (`-aoff`) and using the multipole algorithm (`-mdirect`). The problem is small enough to fit in memory with all the matrices formed directly however is large enough to warrant using the iterative algorithm instead of LU decomposition.

For the  $21 \times 25$  problem, the time spent forming the preconditioner is significant. The `-pdiag` option could have been used to reduce the preconditioner time, however this would come at the expense of slower convergence especially at the higher frequencies.

### 3 Geometry Postscript Pictures

The ability to see the geometry under analysis is an important tool for debugging input files. This section describes how to generate postscript files for visualizing the three dimensional structures defined by a given input file. The process involves first running FastHenry to generate a file of “panels” and then running `zbuf` to generate the postscript file.

#### 3.1 Creating the panel file

When FastHenry is given the `-f` option, it changes from calculation to visualization mode. It will produce a file of panels readable by the `zbuf` program. Each panel is a quadrilateral representing one of the faces of a segment in the input geometry. This file of panels can either be formed from the segments of the geometry before FastHenry’s refinement with the `simple` argument, or after refinement with the `refined` argument. The basenames for the files produced are `zbuffile` for simple geometry and `zbuffile2` for refined unless the `-S` is used.

In addition, the file `zbuffile_shadings` or `zbuffile2_shadings` is produced which specifies a shade of gray to assign to each of the panels. `zbuffile_shadings` is produced for `-f simple` and shades each of the reference planes (if drawn with `-g on`) of the structure differently, and leaves all other structures white. `zbuffile2_shadings`, produced with `-f refined`, does the same, however it attempts to shade differently each of the conductors specified with a `.external` statement.

### 3.2 Creating the postscript file

The panel file described in the previous section can be rendered in postscript with the `zbuf` program. This program is a modified version of the algorithms used for visualization in the capacitance extraction program, FastCap, developed by Keith Nabors.

The complete command line format of the `zbuf` program is

```
zbuf [-a<azimuth>] [-e<elevation>] [<input-file>]
      [-r<rotation>] [-h<distance>] [-s<scale>] [-w<linewidth>]
      [-u<upaxis>] [-q] [-x<axeslength>]
      [-b<.figfile>] [-c] [-v] [-n] [-f] [-g] [-m]
```

Table 2 itemizes the options and Table 1 gives several examples of their use.

Usually the default settings produce an acceptable plot, but many adjustments are possible, the most important being the view angles (relative to a coordinate system parallel to the input coordinates and centered on the center of the object) which are adjusted using the `-a` and `-e` options. Other options control view distance (`-h`), two-dimensional projection scale (`-s`), rotation (`-r` and `-u`) and line width (`-w`). Axes of any length may be included with the `-x` option and lines, arrows and dots may be added using `-b`. Shading can be accomplished with the `-q` option. Table 1 gives the commands used to produce the line drawings in this guide as examples.

In version 3.0, the `zbuf` program now takes the “`-m`” argument to produce a Matlab file for faster visualization in matlab. This is very beneficial for large files since producing the postscript file with `zbuf` can take  $n^2$  time. The matlab file can be viewed within matlab with the `fasthenry-3.0/bin/plotfastH.m` matlab function. The file `zbuffile.mat` would be produced with “`zbuf -m zbuffile`” which can then be viewed in matlab with “`plotfastH('zbuffile.mat')`”. Also, you can modify the file `src/zbuf/dump_struct.c` to output in YOUR own format instead of matlab.

### 3.3 Tricks

For Figure 7, the shading of the trace and ground plane were swapped by negating all the shading values with the `awk` file `invert.awk` provided in the `zbuf` source directory. Also, Figure 10, since it is only one pin, would normally appear white and all panels would have shade 0 in the shading file. By randomly changing one panel to `-5` and another to `10`, then `-5` becomes white, `10` becomes black, and `0`, a shade of gray. The single black panel is seen on the left end of the figure and the white panel is obscured by other panels.

### 3.4 Visualization artifacts

This section describes possible problems with producing pictures. Many of the listed items come directly from the FastCap documentation.

1. To minimize the number of panels that `zbuf` must deal with, FastHenry does not output panels for the end faces of a segment. Thus, segments often appear hollow. Uncomment the appropriate regions in `writefastcap.c` to change this.

Figure	FastHenry/Zbuf Usage
5	<code>fasthenry gpexamp_copper.inp -f simple -S _pcb zbuf zbuffile_pcb</code>
6	<code>fasthenry vias.inp -f refined -g on zbuf -a60 -e60 zbuffile2 -q</code>
7	<code>fasthenry holey_gp.inp -f simple -g on mv zbuffile holey_gp awk -f src/zbuf/invert.awk zbuffile_shadings &gt; holey_gp_shadings zbuf holey_gp -e0 -q -uy</code>
8	<code>fasthenry pin-connect.inp -f simple -S pin zbuf zbuffilepin -a25</code>
9	<code>fasthenry 30pin -f refined zbuf zbuffile -a30 -q</code>
10	<code>fasthenry one_of_30pin -f simple zbuf zbuffile -q -e20 -a0</code>

Table 1: Commands used to generate some representative figures in this guide.

zbuf Options			
Option	Default	Range	Function
-a	50.0	†	Specifies azimuth view angle in degrees.
-e	50.0	†	Specifies elevation view angle in degrees.
-r	0.0	†	Specifies final rotation of 2-D image in degrees.
-h	2.0	$\geq 0.0$	Specifies distance from surface of object in object radii.
-s	1.0	$> 0.0$	Specifies final scaling of 2-D image.
-w	1.0	$\geq 0.0$	Specifies postscript file line width.
-u	z	x, y, or z	Specifies which 3-D axis is mapped to y-axis in 2-D image.
-q	—	—	Uses <code>zbuffile_shading</code> shadings file to give grayscale shading to conductors.
-x	1.0	$> 0.0$	Includes axes of length <code>axeslength</code> in picture.
-b	◇	◇	Specifies lines, dots and arrows to superimpose on picture.
-c	—	—	Puts the command line in postscript file pictures.
-v	—	—	Removes <code>showpage</code> from postscript file.
-n	—	—	Numbers faces in order input.
-f	—	—	Suppresses hidden line removal.
-g	—	—	Prints the graph used to order the panels in postscript file.
-m	—	—	Produces a matlab file rather than a postscript file

Table 2: Zbuf Options

†Range is unrestricted.

◇ See function `readLines()` in `src/zbuf/zbufInOut.c` for a description of the `.fig` file format.

2. `zbuf` sometimes produces panel ordering errors (and usually print warning messages) when panels intersect. However, often the postscript file is correct.
3. Occasionally the postscript pictures are incorrect even for legal discretizations. There are two known bugs leading to this problem. The first is often avoided by changing the view angle slightly. The second is caused by problems with dimensions on the order of  $10^{-4}$  or less. This problem can only be avoided by rescaling the input by changing the `.Units` to `km`.
4. The panel sorting algorithm used to write out the postscript files takes time proportional to the number of panels squared. This limits its usefulness to problems with fewer than a few thousand panels.
5. The shading algorithm uses one shade of gray to shade all the segments along only one path between the two nodes of a `.external` statement. Thus, if there are multiple paths, such as in Figure 9, the other paths will appear white.
6. When axes are included in postscript files using the `-x<axes length>` option, the axes' two-dimensional projections appear in the postscript file before the panel fills. This means that the object should be between the view point and the axes, otherwise the axes can be obscured strangely. Thus the option works best when viewing objects that lie entirely in the positive orthant from a view point in the positive orthant.

## 4 Reference Plane Current Visualization

When FastHenry is given the `-d grids` option, it dumps Matlab readable files of the current distribution in each reference plane of the geometry under analysis. One file is dumped for each of the columns computed for the admittance matrix. The file holds the current distribution produced by setting the source across the port corresponding to the current column to one volt and all others ports to zero. This process is then repeated for each frequency point.

**Note:** For nonuniformly discretized planes, the filename is the same, but the file format is very different. See `nonuniform_manual.ps`. This document also describes viewing the current distribution in other than reference planes.

### 4.1 Current Files

The files are named according to the column in `Zc.mat` and frequency point. They have the form `Gridn_m.mat` where  $n$  is the column (and row) in `Zc.mat` of the port which has the one volt source and  $m$  is the frequency point number. The frequency points start at 0 and increment by one for each new frequency.

Inside each file are matrices for each reference plane named `grid1g<name>` and `grid2g<name>` where *name* is the name of the reference plane defined in the input file. Each entry in `grid1g<name>` is the complex valued current phasor for a filament in the `seg1` direction. `grid2g<name>` corresponds to filament currents in the `seg2` direction.

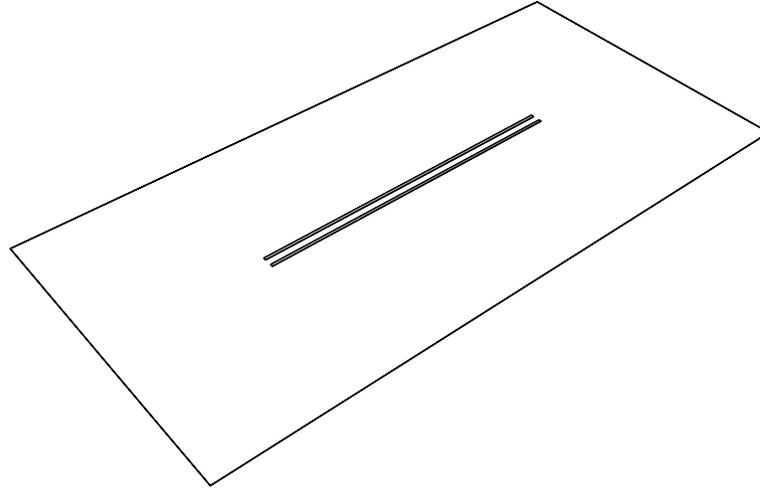


Figure 14: Two traces over a ground plane

In general it is difficult to visualize complex current, but when the current is purely real (DC) or purely imaginary (high frequency), these two matrices can be thought of as the current in two orthogonal directions.

## 4.2 Examples

### 4.2.1 Traces over a solid plane

Consider the example file `together.inp` shown in Figure 14. This structure is simply two traces passing over a ground plane with one end of each trace shorted to the plane. Thus, when one volt is applied to the other end of one of the traces, current travels down the trace and returns through the plane. To generate the current distribution files:

```
fasthenry together.inp -d grids -x trace1
```

The files `Grid1_0.mat` and `Grid1_1.mat` are produced in addition to FastHenry's normal output files. By default, FastHenry would compute the current distribution for the source voltage placed one trace and then again for the other trace. Since the current distribution is nearly identical for both cases, we need only compute one of these cases and thus the `-x` option is used to specify that only the column of the admittance matrix corresponding to port `trace1` is to be computed (see input file). Also, since visualization is the goal, `together.inp` specifies that only a low frequency case,  $f = 10^{-1} Hz$ , and a high frequency case,  $f = 10^{19} Hz$  are to be computed. In this example, the discretization of the plane is coarse for visualization purposes.

The current distribution can then be viewed using the following Matlab commands to produce Figures 15 and 16.