SANDIA REPORT

SAND91-0516 • UC-274

Unlimited Release

Printed May 1991

SEP 12 1991

PC-1D Installation Manual and User's Guide Version 3.1

Paul A. Basore

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550
for the United States Department of Energy
under Contract DE-AC04-76DP00789

SF2900Q(8-81)

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
Office of Scientific and Technical Information
PO Box 62
Oak Ridge, TN 37831

Prices available from (615) 576-8401, FTS 626-8401

Available to the public from
National Technical Information Service
US Department of Commerce
5285 Port Royal Rd
Springfield, VA 22161

NTIS price codes
Printed copy: A05
Microfiche copy: A01

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency Thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

Distribution Category
UC-274

Unlimited Release
SAND91-0516
Printed September 1991

SAND--91-0516
DE91 018307

PC-1D

INSTALLATION MANUAL AND USER'S GUIDE

Version 3.1

Paul A. Basore
Photovoltaic Technology Research Division
Sandia National Laboratories
Albuquerque, NM 87185

Abstract

PC-1D is a software package for personal computers that uses finite-element analysis to solve the fully-coupled two-carrier semiconductor transport equations in one dimension. This program is particularly useful for analyzing the performance of optoelectronic devices such as solar cells, but can be applied to any bipolar device whose carrier flows are primarily one-dimensional. This User's Guide provides the information necessary to install PC-1D, define a problem for solution, solve the problem, and examine the results. Example problems are presented which illustrate these steps. The physical models and numerical methods utilized are presented in detail. This document supports version 3.1 of PC-1D, which incorporates faster numerical algorithms with better convergence properties than previous versions of the program.

MASTER

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

for

CONTENTS

	<u>Page</u>
1. INTRODUCTION	1
2. INSTALLATION	3
2.1 System Requirements -----	3
2.2 Diskette System -----	4
2.3 Fixed Disk System -----	7
3. INPUT PARAMETERS	11
3.01 Proceed with Solution -----	13
3.02 Preview Parameters/Examine Solution -----	13
3.03 Simulation Parameters -----	13
3.04 Output Graphs -----	14
3.05 Reinitialize -----	14
3.06 Number of Finite Elements -----	15
3.07 Solution Mode -----	16
3.08 Temperature -----	17
3.09 Area -----	17
3.10 Thickness -----	18
3.11 Material -----	19
3.12 Doping -----	24
3.13 Recombination -----	25
3.14 Surfaces -----	27
3.15 Circuit -----	27
3.16 Collector/Base Voltage -----	29
3.17 Light -----	29
3.18 Photogeneration -----	30
3.19 Reflectance -----	31
4. PROBLEM SOLUTION	35
4.1 Solution Phases -----	35
4.2 Examination Mode -----	36
4.3 User-Defined Graphs -----	40
4.4 Batch Mode -----	41
4.5 Convergence Failure -----	44
5. NUMERICAL METHOD	47
5.1 Fundamental Equations -----	47
5.2 Discretization -----	51
5.3 System Solving -----	54
5.4 Node Placement -----	57
6. EXAMPLE PROBLEMS	59
6.1 MOS Inversion Layer (MOS.PRM) -----	59
6.2 Forward-Biased Diode (DIODE.PRM) -----	61
6.3 Solar Cell (CELL.PRM) -----	66
6.4 Heterojunction Transistor (HBJT.PRM) -----	73
7. PC1DPLOT	77
8. BIBLIOGRAPHY	79

FIGURES

	<u>Page</u>
1. PC-1D program control menu	12
2. PC-1D front-surface facet model	18
3. PC-1D external circuit elements	28
4. Default four-graph display	37
5. Examination mode parametric graph	39
6. User-defined graph menu	41
7. MOS inversion layer charge density	60
8. Diode doping profile	62
9. Equilibrium diode electric field	63
10. Forward-biased diode carrier concentrations	64
11. Diode turn-off transient circuit	65
12. Diode turn-off transient	67
13. Solar cell I-V circuit with curve-smoothing	68
14. Solar cell I-V curve	69
15. Solar cell quantum efficiency versus wavelength	71
16. Gummel plot displaying transistor gain	74
17. Common-emitter characteristic curve	75

TABLES

	<u>Page</u>
1. Commonly Encountered Indices of Refraction	32
2. Solution Phases	35
3. Available Spatial Functions	42
4. Available Temporal Functions	43
5. PC1DPLOT Data File Format	78

PC-1D will turn your personal computer into a powerful workstation for investigating the transport of electrons and holes in semiconductor devices. You will soon discover that PC-1D's highly interactive graphical environment and computational efficiency more than make up for the modest speed of today's personal computers. Whether your interest is in teaching or research, PC-1D's pop-up windows and flexible internal models make it easy to set up and modify a broad range of one-dimensional problems from MOS capacitors to bipolar transistors. Heterojunctions, arbitrary doping profiles, doping-dependent mobility, bandgap narrowing, Auger recombination, interface recombination, temperature effects, photogeneration of carriers, complex boundary conditions, and transient solutions are some of the powerful features built into this program.

PC-1D runs on the IBM PC or PS/2 and compatible microcomputers which are equipped with an 8087-family numeric data coprocessor, graphics, and at least 512 Kbytes of RAM. A more detailed description of the system requirements is given in Section 2.1. If you are in a hurry to see what PC-1D can do, begin with the installation instructions in Section 2, then skip to the example problems in Section 6. Later, you'll want to come back and read Sections 3 and 4 to understand how to set up and run your own problems. Section 5 is for users who need an in-depth understanding of the models and limitations of PC-1D. Section 7 describes in detail the format used for storing and plotting data using the auxiliary program PC1DPLOT. Section 8 contains a bibliography of relevant material. If, after reading this User's Guide, you have problems either installing or running PC-1D, you may call ISURF at the number shown on the cover page. They will direct you to one of the program authors for assistance.

The development of this program was initiated by Dr. Paul A. Basore while he was a member of the faculty of electrical engineering at Iowa State University in 1984, with funding from the photovoltaic technology division of Sandia National Laboratories. The project was originally intended to take advantage of the interactive environment of personal computers to help analyze solar cells. Additional funding from IBM Corporation in the form of a Faculty Development Award made it possible to expand the program into a full-featured semiconductor analysis package for public distribution. Version 1 was released in December 1985. Key contributors to the first version included students Diane T. Rover, Gregory M. Thorson, and Gregory D. Anderson. After moving to Sandia in 1986, Basore began a comprehensive rewrite of the program, with assistance from Barry R. Hansen and graduate student Arlynn W. Smith of the Georgia Institute of Technology. Version 2 was first demonstrated publicly in September 1988 and released in March 1989. Version 3 is fully compatible with Version 2, but runs faster and converges to a solution more reliably. You can now create a file that lists all of the parameters used for a given solution, and you can create auxiliary space-time graphs. Version 3 was released through the Iowa State University Research Foundation in April 1991.

2

INSTALLATION

2.1 System Requirements

The minimum system configuration for PC-1D requires:

1. An IBM PC, PC/XT, PC/AT, PS/2, or compatible,
2. PC-DOS 2.1 or later,
3. 512 Kbytes of RAM,
4. An 8087-family numeric coprocessor chip,
5. The IBM CGA, EGA, or VGA graphics adapter and display, and
6. One diskette drive (360 Kbytes or more).

Optional features supported include:

1. Mouse: Microsoft, Mouse Systems, IMSI, TopView, PS/2, or compatible,
2. Pen Plotter: IBM, HP, or compatible.

640 Kbytes of installed RAM is necessary when either EGA or VGA graphics modes are used, or if a plotter is to be used. PC-DOS 3.0 or later is required to operate the program in batch mode (Section 4.4).

Even with 640 Kbytes, you may not have enough RAM to run PC-1D in its full configuration if you have other memory-resident programs installed. If you don't have enough free RAM available, PC-1D may not load at all. If it does, it may still give you a "No Room in Heap" error message if it runs out of memory during the solution of a problem. You can reduce the memory requirements of PC-1D by using a lower-resolution display driver (EGA instead of VGA, CGA instead of EGA) and by not loading the plotter driver.

This program produces video output in graphics mode using the IBM Virtual Device Interface. Graphic display drivers are provided on the distribution diskette for the following configurations:

1. CGA Color Graphics Adapter and Display (2-color, 640x200)
2. EGA Enhanced Graphics Adapter and Display (16-color, 640x350)
3. VGA Virtual Graphics Array and Display (16-color, 640x480)

Other video adapters and displays that support a horizontal resolution of at least 640 pixels (such as Hercules) can be used if a corresponding VDI display driver can be obtained. You will need to specify that you want a driver that is compatible with the IBM Graphics Development Toolkit Version 1.2. Drivers for a variety of systems are available from:

Graphical Software Systems, Inc.
9590 SW Gemini Drive
P.O. Box 4900
Beaverton, OR 97005
(503) 641-2200

Several of the older "IBM-compatible" computers lack one important aspect of IBM-compatibility for running PC-1D. These computers fail to set a specific memory location that informs PC-1D (and other programs) whether an 8087-family numeric coprocessor is installed. If you get an error message that says you don't have a coprocessor, when you are certain that you do, try running the program 8087ON.COM before you run PC-1D. This program is provided on the PROGRAM distribution diskette.

Note to Users Upgrading from PC-1D Version 1 or 2:

PC-1D/3 is completely file compatible with PC-1D/2 and uses the same VDI device drivers. Any material or parameter files you may have developed can be used without modification. The only thing you need to do is update the files PC-1D.EXE and PC1DPLOT.EXE. You may wish to update the file SI.MAT if you want to use the new default material parameters for silicon. The VDI drivers for PC-1D/3 are not compatible with those for PC-1D/1. You cannot run both programs on the same computer unless you modify CONFIG.SYS and reboot the machine in between. The following instructions assume that all PC-1D/1 files (including MRES.SYS and HRES.SYS), have been removed prior to installing PC-1D/3.

2.2 Diskette System

If you plan to install PC-1D on a fixed disk, skip this section and proceed directly to Section 2.3, entitled "Fixed Disk System."

PC-1D cannot be run directly from the distribution disks. You must copy certain files from the distribution disks onto your own formatted diskettes in order to create the necessary environment for the program. The instructions which follow will produce three diskettes. One will set up your system properly for PC-1D when you boot your computer from that diskette. The second diskette will store data, and the third diskette will contain the program itself.

Start with three freshly-formatted diskettes (360 Kbytes or more). Your DOS manual has instructions for formatting disks. One of these disks, which will be called the "boot disk," should contain the PC-DOS (MS-DOS) operating system, version 2.1 or later. You will need version 3.0 or later if you want to run the program in batch mode (Section 4.4). Do not put the operating system on the other two disks, which will be called the "data disk" and the "program disk".

First, prepare the boot disk. You will need to create a file which has special meaning when the system is booted: CONFIG.SYS. This is a simple text file, the contents of which depends on what type of graphics adapter and display you have, and whether you have a mouse or plotter. Several files that have the name CONFIG with various extensions are provided on the distribution disk labeled PROGRAM. You will copy one or more of these

files onto your boot disk to create a single file with the name CONFIG.SYS. Follow only the instructions below which correspond to your hardware configuration. The examples given in bold type assume that the PROGRAM distribution disk is in drive B: and your formatted boot disk is in drive A:.. Additional instructions on how to copy files can be found in your DOS manual.

For ALL hardware configurations:

COPY B:CONFIG.SYS A:

If you have the IBM Color Graphics Adapter:

COPY A:CONFIG.SYS+B:CONFIG.CGA A:

If you have the IBM Enhanced Graphics Adapter with EGA Color Display:

COPY A:CONFIG.SYS+B:CONFIG.EGA A:

If you have the IBM Virtual Graphics Array with VGA Color Display:

COPY A:CONFIG.SYS+B:CONFIG.VGA A:

If you have the Microsoft Serial Mouse:

COPY A:CONFIG.SYS+B:CONFIG.MSS A:

If you have the Microsoft Parallel Mouse:

COPY A:CONFIG.SYS+B:CONFIG.MSP A:

If you have the Mouse Systems PC Mouse or IMSI Optimouse:

COPY A:CONFIG.SYS+B:CONFIG.PCM A:

If you have the Visi-On Serial Mouse:

COPY A:CONFIG.SYS+B:CONFIG.VIS A:

If you have the Top-View Logical Mouse:

COPY A:CONFIG.SYS+B:CONFIG.TVM A:

If you have the IBM PS/2 Mouse:

COPY A:CONFIG.SYS+B:CONFIG.MOU A:

If you have an IBM 7371, 7372, 6180, or HP pen plotter:

COPY A:CONFIG.SYS+B:CONFIG.PLT A:

For ALL hardware configurations:

COPY A:CONFIG.SYS+B:CONFIG.VDI A:

COPY B:AUTOEXEC.BAT A:

COPY B:INIT_VDI.EXE A:

Note that the line for VDI.SYS must come after all other VDI file specifications in your CONFIG.SYS file.

Create a subdirectory on your boot disk called PC1D to hold all of the device drivers. Do this by issuing the following DOS command:

MD A:\PC1D

Now remove the PROGRAM distribution diskette from drive B: and replace it with the DRIVERS distribution diskette. For ALL hardware configurations:

COPY B:*.SYS A:\PC1D

If your mouse or plotter is connected to a serial port and that port is not COM1, you will need to use your text editor to insert lines similar to the following at the beginning of the AUTOEXEC.BAT file on your boot disk:

set vdigimou=com2

set vdiplotr=com3

Display the contents of your boot disk's CONFIG.SYS file using

TYPE A:CONFIG.SYS

to determine which device driver name to use in place of vdigimou for the mouse in the above example. These commands must come before the line in AUTOEXEC.BAT that invokes the program INIT_VDI.EXE.

If you are planning to use the plotter, you must also set up the serial port to which it is connected. Use the MODE command supplied with DOS to accomplish this. You will have to copy the DOS file MODE.COM from your DOS system disk to your boot disk. Then you must add a line in your AUTOEXEC.BAT file on your boot disk, before the line that invokes the program INIT_VDI.EXE. This line should be similar to the following:

mode com1:9600,n,8,1,p

Check your DOS manual for information on how to select the parameters for the MODE command.

This completes the preparation of the boot disk. You may wish to write-protect this disk to prevent accidental erasure.

To prepare the program disk, remove your boot disk from drive A: and replace it with your empty, formatted, program disk. Also, put the PROGRAM distribution diskette in drive B:. Then copy just one (albeit very large) file:

COPY B:PC-1D.EXE A:

That completes the preparation of the program disk. You will probably want to write-protect this disk to prevent accidental erasure.

To prepare the data disk, remove the program disk from drive A: and replace it with your empty, formatted, data disk. Also, remove the PROGRAM distribution disk from drive B: and insert the DRIVERS distribution disk in drive B:. Copy all of the files on the distribution disk to your data disk:

COPY B:*. * A:

Some of the files you just copied to your data disk were already copied to your boot disk, and are not needed on the data disk. You should delete these files from your data disk using this command:

DELETE A:*.SYS

This completes the preparation of the data disk. You will be using this same disk to save data generated by PC-1D, so do not write-protect this disk.

To execute the program, first boot your system from the **boot disk** in drive A: using Ctrl-Alt-Del. This will load the correct CONFIG.SYS environment and execute the commands in your AUTOEXEC.BAT file. Invoke the program by replacing the boot disk with the **program disk** and typing PC-1D in response to the DOS prompt. The system will give you a warning message saying that it cannot find the data files. Replace the program disk with the **data disk**, then press any key to start the program. Leave the data disk in place as long as PC-1D is running. After you have exited PC-1D, you can plot data that you saved while running PC-1D by typing PC1DPLOT in response to the DOS prompt (Section 7). If you should want to run PC1DPLOT at some later time without running PC-1D first, you should boot the system using the boot disk, then insert the data disk and run PC1DPLOT.

2.3 Fixed Disk System

This section is only for users who wish to install PC-1D on a fixed disk. If you have a diskette system, you should have installed the program using the directions in Section 2.2, and may skip the material in this section.

PC-1D is a large program, and will load much faster when installed on a fixed disk. Once the program starts, however, it will run at the same speed whether loaded from a fixed disk or a diskette. You might want to consider setting up PC-1D for diskettes, even if you do have a fixed disk, to avoid loading the device drivers every time you turn on your computer.

In order to establish the correct operating environment for PC-1D, you must copy several files from the distribution disks to your fixed disk. You will add lines to your CONFIG.SYS and AUTOEXEC.BAT files in the root directory, and transfer the remaining files into a subdirectory labeled PC1D.

Place the distribution disk labeled PROGRAM into your disk drive. The following instructions assume that this is drive A:. If not, just substitute the appropriate drive letter for A: in the following directions. Make certain that you are in the root directory of drive C: by typing:

C> C:

```
C> CD \
```

You are likely to already have a CONFIG.SYS file in the root directory of your fixed disk. This file controls the loading of device drivers which are specific to your hardware configuration. If you do not have a CONFIG.SYS file, then copy the blank CONFIG.SYS file from the distribution disk:

```
C> COPY A:CONFIG.SYS C:
```

Proceed to add lines to your CONFIG.SYS file using the following commands which correspond to your hardware configuration. Enter only those commands that correspond to your computer system. Additional information on copying and appending files is contained in your DOS manual.

If you have the IBM Color Graphics Adapter:

```
COPY C:CONFIG.SYS+A:CONFIG.CGA C:
```

If you have the IBM Enhanced Graphics Adapter with EGA Color Display:

```
COPY C:CONFIG.SYS+A:CONFIG.EGA C:
```

If you have the IBM Virtual Graphics Array with VGA Color Display:

```
COPY C:CONFIG.SYS+A:CONFIG.VGA C:
```

If you have the Microsoft Serial Mouse:

```
COPY C:CONFIG.SYS+A:CONFIG.MSS C:
```

If you have the Microsoft Parallel Mouse:

```
COPY C:CONFIG.SYS+A:CONFIG.MSP C:
```

If you have the Mouse Systems PC Mouse or IMSI Optimouse:

```
COPY C:CONFIG.SYS+A:CONFIG.PCM C:
```

If you have the Visi-On Serial Mouse:

```
COPY C:CONFIG.SYS+A:CONFIG.VIS C:
```

If you have the TopView Logical Mouse:

```
COPY C:CONFIG.SYS+A:CONFIG.TVM C:
```

If you have the IBM PS/2 Mouse:

```
COPY C:CONFIG.SYS+A:CONFIG.MOU C:
```

If you have an IBM 7372, 7371, 6180, or HP pen plotter:

```
COPY C:CONFIG.SYS+A:CONFIG.PLT C:
```

For ALL hardware configurations:

```
COPY C:CONFIG.SYS+A:CONFIG.VDI C:
```

Note that the line for VDI.SYS must come after all other VDI file specifications in your CONFIG.SYS file.

If your mouse or plotter is connected to a serial port and that port is not COM1, you will need to use your text editor to add lines to your AUTOEXEC.BAT file in your root directory similar to the following:

```
set vdigimou=com2
```

```
set vdiplotr=com3
```

Type out the contents of your CONFIG.SYS file using

```
TYPE CONFIG.SYS
```

to determine which device driver name to use in place of vdigimou for the mouse in the above example. If you don't already have an AUTOEXEC.BAT file, create one using the directions included in your DOS manual. Note that addition of these lines may lead to a DOS error message "Out of Environment Space." If you get this error message, type **SET** in response to the DOS prompt. You will get a listing of your environment. The number of characters that are allowed in the environment is very limited. Many people use up all of their environment space with lengthy PATH commands. Note that you only need to include a directory in the PATH command if you want to invoke programs located in that directory from other directories. You probably only need to include your C:\DOS subdirectory, the directory for your text processor, and perhaps any programming-language directories you may have. See your DOS manual for more information about what belongs in the PATH command.

If you are planning to use the plotter, you must also set up the serial port to which it is connected. Use the MODE command supplied with DOS to accomplish this. If your DOS files are in a subdirectory, make certain that the directory that contains your DOS system files (including MODE.COM) is included in the PATH command in your AUTOEXEC.BAT file. Then you must add a line in your AUTOEXEC.BAT file on your boot disk, after the PATH command. This line should be similar to the following:

```
mode com1:9600,n,8,1,p
```

Check your DOS manual for information on how to select the parameters for the MODE command.

Create a directory called PC1D so that all of your PC-1D files can be grouped together. Do this from within the root directory by issuing the following DOS commands:

```
C> MD PC1D
```

```
C> CD PC1D
```

The second line changes the active directory to PC1D. If you want to use a name other than PC1D, you will have to modify your CONFIG.SYS file accordingly. Copy one (albeit very large) file from the PROGRAM disk into the PC1D directory.

```
C> COPY A:PC-1D.EXE C:
```

Remove the PROGRAM distribution disk from drive A: and replace it with the distribution disk labeled DRIVERS. Copy all of the files from the DRIVERS disk into the PC1D directory.

COPY A:*. * C:

To execute the program, first remove the disk from drive A: and reboot your system using Ctrl-Alt-Del. This will load the correct CONFIG.SYS environment and execute the commands in your AUTOEXEC.BAT file. Invoke the program itself by changing to the PC1D directory and typing PC-1D in response to the DOS prompt. You can later plot data that you have saved while running PC-1D by instead typing PC1DPLOT in response to the DOS prompt (Section 7).

There are two different ways to invoke PC-1D. If you simply type PC-1D in response to the DOS prompt, the program will read the file PC-1D.DAT to determine which set of parameters were used the last time you ran the program, and will load those parameters. The parameters are stored in a binary-encoded file with a PRM suffix (e.g. MOS.PRM). Each of these parameter files includes all of the on-screen parameters to describe the device to be analyzed and what kinds of excitation (voltages, light) to apply. The parameter file also specifies external data files that may be necessary to describe parameters that are a function of position, wavelength, or time; but these files are not loaded by the program until it is time to begin solving the problem. These external data files are described in more detail in the following subsections. Provided there is enough RAM memory to load the program, and all necessary device drivers have been setup properly in CONFIG.SYS, the program will present you with an introductory screen and commence with interactive input.

If you have DOS 3.0 or later, you can invoke the program in batch mode by typing

PC-1D DIODE DATA1 DATA2 DATA3 DATA4

in response to the DOS prompt. The addition of an input parameter filename (DIODE) and from one to four additional output filenames (DATA1, etc.) notifies the program that batch mode is desired. In this case, the problem specified by DIODE is solved without any further interaction with the user, and the results are stored in the files DATA1.PDF, DATA2.PDF, DATA3.PDF, and DATA4.PDF. Batch mode is described in more detail in Section 4.4. The remainder of this section assumes that you did not invoke the program in batch mode, but instead have entered the interactive-input mode. Indeed, you have to run the program in interactive mode first, to set up the parameter files that you can later submit for batch-mode solution.

PC-1D uses pop-up menus for interactive input. The structure of these menus has been designed to allow for rapid examination and alteration of a large number of parameters. Within a particular menu you can either move the mouse or use the cursor-arrow keys to maneuver the cursor from item to item. The PgUp and PgDn keys on the keyboard can also be used to jump to the top or bottom item in the menu. The only other keys that are active are Enter and Esc. Use these to either select an item from the menu or to exit the menu, respectively. The left mouse button plays the same role as the Enter key on the keyboard, while the right button plays the same role as the Esc key on the keyboard. Just remember that you always use Enter (the left button) to "select", and use Esc (the right button) to "exit". If you want to enter a new value for a particular parameter listed in the menu, you have to move the cursor to that item and hit Enter first. You will then be presented with a small input window in the lower-right corner of the screen. You can type the new value, using the backspace key to make

corrections if necessary. Then hit Enter to register the new value and return to the menu. Users familiar with PC-1D Version 1.1 will recognize that this is a significant departure from the spreadsheet-style input used previously, but the new technique is necessary to accommodate the complex device structures permitted in Version 2.

When you first start interactive input, the cursor will be positioned at the top of a menu labeled "PC-1D Program Control." This large menu occupies the left half of the screen, and will remain there until you either begin an interactive solution by selecting the first item in the menu (Section 3.01) or else exit PC-1D by hitting Esc from anywhere in the menu. Don't worry about losing data because you hit Esc accidentally; you will be presented with an exit menu that will ask you to confirm your desire to exit to DOS, with a warning message if you have modified but not saved the current set of parameters.

The program control menu, illustrated in Fig. 1, contains 20 items which describe the device and the applied excitation, although the last two lines are not displayed if the light intensity is set to zero. Most of the items in this menu, when selected, lead to additional menus in which the individual parameters that specify the problem to be solved can be examined and modified.

PC-1D PROGRAM CONTROL
Proceed with Solution
Preview / Examine (None)
Simulation Parameters: CELL
Output Graphs: Default
Reinitialize: Yes
Number of Finite Elements: 150
Solution Mode: Transient (16 steps)
Temperature: 25 C
Area: Uniform 1 cm ²
Thickness: 300 um
Material: Si
Doping: Internal
Recombination: Internal
Surfaces: Neutral, Neutral
Circuit: Voc Rcc Rc
Collector Voltage: -999 mV to 999 mV
Base Voltage: Zero
Light: 100 mW/cm ²
Photogeneration: Spectrum AM15G
Reflectance: Coated, 0.85 (Tex)

Fig. 1 PC-1D program control menu.

The subsections which follow describe the input parameters associated with each line item in the Program Control menu. The example problems in Section 6 are an excellent way to practice setting up and modifying parameters using the interactive input mode.

3.01 Proceed with Solution

Select this item to begin solving the device equations for the problem described by the current set of parameters. PC-1D will load any external files that are specified to describe functions of position, wavelength, or time. If any of these files cannot be found or contain format errors, the solution is aborted and you are returned to interactive input. If the problem is set up successfully, the screen will go blank except for a narrow information bar at the bottom of the screen until the initial solution has been completed. Section 4.1 describes the steps that are taken to arrive at the final solution.

3.02 Preview Parameters / Examine Results

On selecting this item you must choose between Preview Parameters and Examine Results. These two options are closely related. The only difference is that Examine Results allows you to review plots of previously-obtained results, whereas Preview Parameters sets up the parameters for the next solution so that you can check to make sure that the doping profile, recombination lifetime, mobility, and cross-sectional area are what you intended. Selection of either item will transfer you to the interactive-examination mode described fully in Section 4.2. The solution state shown on the Preview / Examine line indicates the last successful solution (None, Equilibrium, Steady State, or Transient). Note that if you escape from a solution while it is still in progress and then return to examine the results, the data you see represents an interim state that is not a valid solution of the semiconductor equations.

3.03 Simulation Parameters

This line shows the filename for the parameters that describe the current problem. This label can be up to eight characters, corresponding to a DOS filename. The file extension PRM is assumed and is not displayed. Select this item if you want to load a new parameter file from the disk, revise the parameters in the current file, create a new file, delete a file, or create a parameter-listing file. If you ask to load a new file, the available parameter files in the current disk directory are listed for you to choose from. The "Revise" option simply updates the file on disk to reflect the current set of parameters in the event that some were modified. If you select "Create a New File", you must provide a filename. PC-1D checks to make sure that you don't already have a file with that same name before creating the new parameter file. The "Delete" option lets you erase one parameter file from the current list.

When a new PRM file is loaded, all of the on-screen solution parameters are replaced with the values that were in effect when the PRM file was created or last revised. It is particularly important to understand the effect this has on the material parameters, which are on-screen values. Suppose, for example, that you performed a simulation using the SI.MAT material data

file and saved a PRM file representing a problem of interest. Suppose further that at some later date you decide to revise the SI.MAT data file to include more up-to-date parameter values (Section 3.11). If you then load the PRM file you had saved previously, the screen will show that the program is using the material data from SI.MAT, but in fact it will be using the values that were in SI.MAT when the PRM file was saved, not the new values. Note that this "memory" feature applies only to on-screen parameters. Parameters that are supplied from external files with suffix DOP, TAU, ARA, GEN, BND, INR, ABS, SPC, REF, VLT, or LGT are read each time a solution is run and therefore will always represent the most recent versions of those parameters.

PRM files are binary-encoded files that cannot be examined using a text editor. Each parameter is represented by a 16-byte binary record. To create a readable ASCII file of the parameters for the current simulation, select the "Create Parameter List File" option. A disk file will be created with the same name as the PRM file but with an LIS suffix. The labels used for each parameter in the listing file are read from the file PARMS.TXT (which can be edited). Use any text editor to print the file.

3.04 Output Graphs

This line shows either "Default" or "Specified", and may additionally show "Auxiliary". The default output graphs are Carrier Concentrations, Electrostatic Potential, Current Density, and Quasi-Fermi Energies. By selecting "Specified", you can pick any four graphs from the list of predefined plots. Returning to "Default" does not erase your selections; you can switch back to "Specified" and your previous choices will be restored. The four graph selections will be displayed at the successful completion of each phase of the solution: equilibrium, steady state, and after each time step for transients. These selections also determine which data will be saved at the conclusion of a batch-mode solution.

If you select Auxiliary, you are given the opportunity to choose any of the available spatial functions (Section 4.3, Table 3), and then asked to specify a distance from the front surface of the device. PC-1D will store the value of the function chosen at the location specified for each time step in a transient solution. The resulting function of time can then be examined by selecting the Auxiliary Space-Time Function (Section 4.2).

3.05 Reinitialize

This line is either Yes or No. When you select this item you simply toggle between the two possible values. This field is initialized to "Yes" every time you enter the interactive input mode. By setting this field to "No", the current solution (node placement, doping levels, electrostatic potential, and carrier concentrations) is used as a starting point for the next solution, rather than solving for equilibrium first. The photogeneration profile from the previous solution may also be used, if no

parameters affecting it have been altered. This can save some time if you know that the previous solution is a good approximation to the next solution. Furthermore, the clock is not reset to $t=0$ if you select "No", so that you can continue with a transient solution that you may have interrupted while it was in progress. You can also use this feature to change the excitation parameters part-way through a transient solution. For example, if after 10 of 20 time steps you change from a fixed voltage of 5 volts to a ramp from 0 to 10 volts, you will get a ramp from 5 to 10 volts for the remaining time steps, since you have already completed half of the total sweep time.

3.06 Number of Finite Elements

This line indicates the number of elements used in the last solution, or the maximum allowed if no solution has been completed. Select this item to examine or modify several parameters that affect the speed, accuracy, and convergence of the numerical solution algorithm. The first of these parameters is the element size factor. This factor is 1.0 by default. This is about as large as this parameter can be set and still obtain reasonable accuracy. Select a smaller number to increase the number of elements used. The specific number of elements depends on the complexity of the device being studied, and may vary significantly from problem to problem. For a given element size factor, however, you can expect to obtain comparable accuracy. The maximum number of elements allowed is 150. If this is not enough to satisfy the element size factor requested, a larger size factor will be used instead, and its value will be displayed. Although a smaller element size factor will increase the accuracy of the solution, it will also increase the time required to obtain the solution. The algorithm used to determine the number and sizing of the elements is described in Section 5.4.

The semiconductor equations are highly nonlinear and require an iterative solution. The solution is said to have converged when the maximum change in any value of electrostatic or quasi-Fermi potential, normalized by the thermal voltage (kT/q), falls below a preset level. This Error Limit can be adjusted, though the default value ($1E-6$) should work in most cases. The convergence of the solution is improved substantially by preventing the potential values from changing too much in a single iteration. The Normalized Potential Clamp limits the change in potential to a certain number of multiples of the thermal voltage. A value of 5 provides adequate convergence, without slowing things down too much. Smaller values may be helpful where convergence is a problem (Section 4.5). The effect of the clamp is most noticeable when applying large reverse bias voltage to a device. A reverse bias of 10 volts, for example, represents about 400 room-temperature thermal voltages. At 5 kT/q per iteration, it will take at least 80 iterations before the solution even begins to converge.

The maximum number of iterations permitted can be specified. This is primarily for setting up problems that will be solved later using batch mode, in which case this parameter insures that the program won't get stuck in an infinite loop. A value of 100 gives ample opportunity for convergence, without wasting too much time on solutions that will never converge. It may be necessary to increase this value if large reverse voltages are involved or the potential clamp is set to a small value. Note that most problems converge after 10 to 20 iterations.

The automatic node-placement algorithm adjusts the size and number of elements for maximum accuracy and rapid convergence. You should almost always request an Equilibrium Renode, unless you have a reason for retaining the initial node positions. By selecting the Steady-State Renode, an adjustment of the nodes can be called for at the beginning of each steady-state solution as well. This is generally advisable, because it provides a better element distribution with relatively little cost in solution time. After each renode operation, the equilibrium solution is repeated to insure that the equilibrium values are fully consistent with the new set of elements. By selecting the Transient Renode, you can also invoke a renode operation at the beginning of each time step. This helps increase accuracy in slow-transient solutions that span a wide range of excitation levels, such as I-V curves. It is not recommended in other situations. The node-placement algorithm is discussed in Section 5.4.

3.07 Solution Mode

This line can be set to Equilibrium, Steady State, or Transient. The choice determines the final state the solution will achieve before stopping for examination of results and returning to interactive input. Unless Reinitialize (Section 3.05) has been set to "No", the solution will always proceed by solving first for Equilibrium, then Steady State, then Transient. If the mode is Transient, this line will also display the total number of time steps requested.

When the Solution Mode is Transient, it is necessary to supply additional details for controlling the time dependence of the solution. The Step Size is the elapsed time between graphed results. Setting this value too small may cause convergence problems. Large values (e.g. 1 second or more) can be used to simulate the quasi-static response of the device to a swept excitation. Examples include current versus voltage and spectral response versus wavelength.

The entry for Number of Steps is the number of time steps of duration Step Size that are desired. Normally, the total duration of the transient will be the product of these two values. It is possible, however, to extend the transient using a different time step size. Set Reinitialize to No (Section 3.05), increase the Number of Steps, and alter the Step Size. The additional steps will use the new time step size. You might do this if the device response involves a rapid initial rise followed by a slow decay. A

maximum of 50 time steps are allowed.

The Separate Solutions per Step entry is provided to permit increased accuracy in certain situations where you can't use smaller time steps. Normally equal to 1, setting this value to 2 or more will cause each time step to, in fact, be solved as a series of short substeps. One situation where you would want large time steps is in constructing a graph showing the evolution of a spatial variable across a wide span of time. Spatial information is only stored for four time steps (Section 4.2), so these would need to be long steps. Many substeps would be necessary in this situation to obtain adequate accuracy.

3.08 Temperature

The temperature of the device is assumed to be uniform. The value is displayed in either kelvin or degrees Celsius. The temperature setting affects the thermal voltage (kT/q) for both the device and the external-circuit diodes, the equilibrium np product, the carrier mobility, the internal optical absorption model, and the occupation of the recombination traps. Other temperature-affected parameters must be modified directly by the user: dopant ionization, recombination trap cross-sections, Auger and band-recombination coefficients, index of refraction, dielectric constant, and surface charge.

3.09 Area

This line indicates whether the cross-sectional area for current flow in the device is Uniform, imported from an external Disk File, or calculated for a device with a Faceted front surface. Although PC-1D solves a one-dimensional approximation to the semiconductor transport equations, it goes one step further by allowing the cross-sectional area, A , to be a gradual function of the solution dimension, x . This "quasi-one-dimensional" solution allows you to obtain approximate solutions for devices that are not strictly one dimensional. A simple example would be a device that involves a hemispherical contact, with current flowing radially away from the contact. For this case, the solution dimension x would be radial, and the area A would increase as x^2 . This parabolic area variation can be entered using an external area file with a filename of the form *.ARA.

External area files are standard ASCII files that contain two numerical values on each line, separated by one or more spaces. The first value on each line is a value of x , in cm. The second value on each line is the cross-sectional area, A , in cm^2 . All values of area must be positive, and the values of x must increase monotonically, starting with $x=0$. The maximum number of lines allowed is 150. The last value of x in this file does not define the thickness of the device, which is determined by a separate entry (Section 3.10). If the range of x in the file is less than the thickness, the final area value is extended for the rest of the thickness. If the range of x exceeds the thickness, only part of the area-

file information is used.

The facet model for cross-sectional area is illustrated in Fig. 2. This feature is intended for devices with etched front surfaces that have a specific facet angle with respect to the plane of the substrate. The solution dimension x is shown in the figure. Note that x is not "straight," and that the cross-sectional area for current flow is larger at $x=0$ than at the back surface by the factor $1/\cos(\alpha)$, where " α " is the facet angle. The facet model assumes that the cross-sectional area decreases with a Gaussian shape from its value at $x=0$ toward the substrate value, with a standard-deviation depth constant equal to one-sixth of the depth of the facets, d . The most important feature of the facet model is that it incorporates the increased carrier recombination that occurs at or near a textured surface as a direct result of the increased surface area. Facet angles of up to 60° are accepted. Note that $\langle 111 \rangle$ facets etched anisotropically into a $\langle 100 \rangle$ -oriented surface make a facet angle of 54.74° . The choice of facet angle and depth also affects the photogeneration profile (Section 3.18).

Fig. 2 PC-1D front-surface facet model.

3.10 Thickness

PC-1D allows you to construct a device having from one to three adjacent "regions." Each region can be a different material, with its own bandgap, mobility, dielectric constant, and optical parameters (Section 3.11). In addition, the doping (Section 3.12) and recombination parameters (Section 3.13) can be specified separately for each region. You can add or remove regions by specifying their thickness. A zero thickness removes a region from the problem; setting it greater than zero adds that region to the

problem. The regions must be added in order, starting with region 1.

The regions are separated by grading layers, the thickness of which can be specified. In a real device, transitions between different materials are never perfectly abrupt. In PC-1D the default grading-layer thickness is 10 nm. Setting this entry to a value that is much smaller causes convergence problems and introduces spikes in the solution that would not exist in the actual device. Larger widths can be specified to simulate an extended grading layer between the material regions. The grading layer is considered to be a part of the region on the right-hand side of the interface. Each nonzero region must therefore be at least as thick as the grading-layer thickness. Specifically, PC-1D will decrease the grading-layer thickness if necessary to insure that it is less than one-half of the narrowest region.

3.11 Material

Up to three filenames are listed next to this entry, corresponding to the disk files that contain the material data parameters for each of the three regions. If a region's thickness is 0, no material filename is shown for that region. The filename suffix of MAT is assumed and is not displayed. If no material data file has been selected, "Internal" is displayed and a set of internal default parameters is used. If any of the material parameters have been altered since the material file was loaded from the disk, an asterisk "*" will appear next to the material filename. The asterisk will disappear when a material data file is loaded from disk, or if the data file is revised to incorporate the modified parameters.

Material data files (*.MAT) are standard ASCII files. They can be examined using any text editor. However, it is not advisable to modify these files except from within PC-1D as described in this section. PC-1D expects these files to have a very specific format, and if this format is corrupted, PC-1D will not be able to read the material data file.

Upon selecting the Material item in the main menu, you first must choose which of the active regions to consider. Only regions with a thickness greater than 0 are active. For each active region you can either examine the current set of material data parameters, or work with the disk files. The disk-file options are to Load a new set of material data parameters, to Revise the current set of material data parameters in order to save any changes that you have made, to Create a new material data file, or to Delete a file. When creating a new file, PC-1D will ask for the new filename and will check to make certain that no file with that name already exists prior to creating the file and saving the material parameters. Material files are supplied with PC-1D for silicon (SI), gallium arsenide (GAAS), indium phosphide (INP), germanium (GE), and several aluminum gallium arsenide alloys ($Al_xGa_{1-x}AS$). You can model other materials simply by modifying one of these files and creating a new file using the new material's name.

Note that it is not necessary to save every variation of material parameters you may wish to explore. When you create a parameter (PRM) file (Section 3.03), the current set of material parameters (including modifications) is saved along with all of the other on-screen parameters that describe your problem. Whenever you call up that PRM file, you will get the modified material parameters. It is only necessary to create a new material file if you plan to use that same set of parameters frequently for other problems.

If you choose to examine the material parameters, you will be presented with a material data menu. The first two lines of this menu lead to menus for setting the parameters for the electron and hole mobility models, respectively. These models set the mobility as a function of the total doping concentration at each position, and as a function of the device temperature. The various coefficients are defined by Eq. (1).

$$\mu = \mu_{\min} \cdot T_n^{\beta_1} + \frac{(\mu_{\max} - \mu_{\min}) T_n^{\beta_2}}{1 + \left\{ \frac{N_D + N_A}{N_{\text{ref}} \cdot T_n^{\beta_3}} \right\}^{\alpha \cdot T_n^{\beta_4}}}, \quad (1)$$

where $T_n = T / 300$ K. The parameters $\mu_{\min}$, $\mu_{\max}$, α , N_{ref} , and temperature coefficients β_1 , β_2 , β_3 , and β_4 are supplied by the user for both electrons and holes. The parameters $\mu_{\min}$, α , and N_{ref} can be specified independently for majority- and minority-carrier conditions. To avoid an abrupt change in mobility at the junction, where the majority and minority carriers switch polarity, the mobility is calculated by combining the majority and minority mobilities into a weighted average, based on the relative concentrations of donors and acceptors at each point. Eq. (1) gives the low-electric-field mobility. A velocity-saturation model is added to properly model transport in high-field regions. The mobility of each carrier is reduced from its low-field value (μ_{LF}) to its high-field value (μ_{HF}) such that its velocity never exceeds the specified saturation velocity, v_{sat} .

$$\mu_{\text{HF}} = \mu_{\text{LF}} / \sqrt{1 + (\mu_{\text{LF}} \cdot E / v_{\text{sat}})^2}. \quad (2)$$

The mobility model does not account for mobility degradation due to carrier-carrier scattering associated with high-level injection.

The index of refraction can be specified using either the Internal model, which assumes that the index is the same for all wavelengths of light, or as a function of wavelength using data imported from an external disk file. External index-of-refraction files are standard ASCII files with a filename of the form *.INR. Each line in the file should contain two numerical values, separated by one or more spaces. The first value on each line is a value of wavelength, in nm. The second value on each line is the index of refraction at that wavelength. The maximum number of lines allowed is 100.

All values of the index must be positive, and the values of wavelength must increase monotonically. If the range of wavelengths provided is less than is required for the solution of a problem, the index of refraction for the first (or last) wavelength is used for all smaller (or larger) wavelengths.

The absorption-file entry displays whether the Internal optical absorption model will be used, or if an external absorption file will be used. The internal model generates the optical absorption coefficient as a function of both wavelength and temperature. An external file can be used to provide accurate absorption information as a function of wavelength for a single temperature. Only absorption that leads to photogenerated carriers should be considered, since PC-1D assumes that every absorbed photon produces a hole-electron pair.

The Internal model for photon absorption includes two direct-gap transitions and eight indirect transitions. The indirect transitions are based on two indirect gaps and either absorption or creation of phonons having two possible energies. Temperature affects the absorption coefficient, α , through a decrease in the energy gaps and through a change in the probability of phonon-assisted transitions. All four energy gaps are assumed to depend on the absolute temperature (T), through a temperature coefficient (β) and temperature offset (γ).

$$E_G = E_{Go} - \frac{\beta T^2}{T + \gamma}. \quad (3)$$

The entries required in the absorption menu are the energy gaps at 300 K. Eq. (3) is used to correct these values to the device temperature.

The contributions to the absorption coefficient due to direct transitions for photons with energy $h\nu$ greater than the energy gap E_{Gdi} are given by

$$\alpha_{di} = A_{di} \sqrt{h\nu - E_{Gdi}}; \quad i=1,2. \quad (4)$$

The eight phonon-assisted transitions for indirect-gap energies E_{Gi} and phonon energies E_{pj} take the form

$$\alpha_{ij}(\pm) = A_{ij} \frac{(h\nu - E_{Gi} \pm E_{pj})^2}{\pm(e^{\pm E_{pj}/kT} - 1)}; \quad i=1,2; j=1,2 \quad (5)$$

The upper signs correspond to phonon-absorbing transitions, while the lower signs are for phonon-creating transitions. Only those terms for which $h\nu \pm E_{pj} > E_{Gi}$ are included. Otherwise, there is no absorption due to that transition. The indirect-gap absorption is clamped for energies greater than the next-higher direct gap. All ten absorption coefficients are added to yield the total absorption coefficient for each incident wavelength of light ($h\nu = hc/\lambda$, where $hc = 1239.8424 \text{ eV}\cdot\text{nm}$).

External absorption files are standard ASCII files with a filename of the form *.ABS. Each line in the file should contain two numerical values, separated by one or more spaces. The first value on each line is a value of wavelength, in nm. The second value on each line is the absorption coefficient, in cm^{-1} . The maximum number of lines allowed is 100. All values of the absorption coefficient must be positive, and the values of wavelength must increase monotonically. If the range of wavelengths provided is less than is required for the solution of a problem, the absorption coefficient for the first (or last) wavelength is used for all smaller (or larger) wavelengths. For this reason, it is important that the last entry in the file give an absorption coefficient of zero. Otherwise, PC-1D will assume that the final (nonzero) value applies even out to very long wavelengths, thereby overestimating the total photogeneration in sunlight. Room-temperature absorption files are provided with PC-1D for several materials (e.g. SI.ABS).

The electrical bandgap narrowing line indicates whether the Internal bandgap narrowing model is to be used, or if this information should be imported from a Disk File. In either case, the equilibrium pn product at each position is increased over the square of the temperature-dependent intrinsic concentration by local bandgap narrowing in the conduction band (ΔE_c) and the valence band (ΔE_v).

$$n_{eq}p_{eq} = n_i^2(T) e^{(\Delta E_c + \Delta E_v)/kT}. \quad (6)$$

The Internal bandgap-narrowing model assumes that the onset of bandgap narrowing occurs when the doping concentration exceeds some threshold amount, and any further doping decreases the effective electrical bandgap in proportion to the natural logarithm of the doping density. All of the bandgap narrowing from donors is assumed to occur in the conduction band, while all that from acceptors is assumed to occur in the valence band.

$$\begin{aligned} \Delta E_c &= C_d \cdot \ln[N_d/N_{do}]; N_d > N_{do}, \\ \Delta E_v &= C_a \cdot \ln[N_a/N_{ao}]; N_a > N_{ao}. \end{aligned} \quad (7)$$

External bandgap-narrowing files are standard ASCII files having a filename of the form *.BND. These files contain three numerical values on each line, separated by one or more spaces. The first value on each line is a value of x, in cm. The second value on each line is the conduction band narrowing (eV), and the third value on each line is the valence band narrowing (eV). The values of x must increase monotonically, starting with x=0. The maximum number of lines allowed is 150. The last value of x in this file does not define the thickness of the device, which is determined by a separate entry (Section 3.10). If the range of x is less than the thickness, the final bandgap narrowing values are extended for the rest of the thickness. If the range of x exceeds the thickness, only part of the bandgap-narrowing file is used. Note that the bandgap narrowing file applies to the entire device, not just one region. If a bandgap narrowing file is selected for any region, that file applies to all regions.

The remaining items in the material parameter menu are adjustable material constants. The intrinsic bandgap and electron affinity are assumed to be uniform throughout each region. N_c/N_v , the ratio of the effective conduction band-edge density of states to the valence band-edge density of states is required to accurately determine the intrinsic energy level within the bandgap. The intrinsic carrier concentration entered should be the value at 300 kelvin. PC-1D will modify this value using the intrinsic bandgap specified to obtain the value of the intrinsic carrier concentration at the specified temperature, using

$$n_i(T) = n_i(T_0) \cdot [T/T_0]^{3/2} \cdot e^{[E_g/2kT_0 - E_g/2kT]}, \quad (8)$$

where $kT_0 = 25.851483$ meV. The value of E_g selected can be chosen to maximize the accuracy of the temperature dependence of this formula, since the value of E_g is not used anywhere else by the program. In silicon, for example, the value of E_g selected for this purpose is somewhat larger than the true material bandgap. The value of n_i calculated from Eq. (8) is then increased exponentially by the local bandgap narrowing described in Eq. (6) to calculate the local equilibrium pn product. The local equilibrium pn product is often referred to as the square of the effective intrinsic carrier concentration, n_{ie} .

$$n_{ie}^2 = n_{eq} \cdot p_{eq}. \quad (9)$$

The dielectric constant requested is the DC permittivity divided by the permittivity of free space ($\epsilon_0 = 8.8541878E-14$ F/cm). The remaining material constants are fundamental recombination parameters that are not significantly affected by doping or material quality. These parameters may be a weak function of temperature, but it is up to the user to modify them accordingly. Auger recombination coefficients can be provided both for light and heavy doping. The total Auger recombination rate is the sum of two terms,

$$R = (C_{nh}N_D + C_{ph}N_A)(pn - n_{ie}^2), \text{ and} \quad (10)$$

$$R = [C_{nl}(n - n_{eq}) + C_{pl}(p - p_{eq})](pn - n_{ie}^2). \quad (11)$$

The first term dominates in heavily-doped regions, while the second term is important in lightly-doped regions that are highly injected. Band-to-band recombination is also included. It has the form

$$R = B(pn - n_{ie}^2). \quad (12)$$

B is replaced by the coefficient B_n for n-type material, and by B_p for p-type material. B_n should equal B_p for pure band-to-band recombination, but separate coefficients adds the flexibility of modeling other recombination mechanisms whose rate is proportional to the pn product, or to approximate the effect of a doping-dependent SRH lifetime.

3.12 Doping

The doping profile for each active region in the device can be generated using Internal models or can be imported from an external Disk File. The internal model allows for a background doping and two diffusions from both the front and back surfaces of each region. The background doping can be specified directly as a doping density (cm^{-3}), or indirectly as a resistivity (Ωcm). If a resistivity is specified, it is converted into a doping density using the current temperature and carrier mobility parameters. The diffusions can have one of four shapes, defined as follows:

Uniform: $N(x) = N_s, x < x_d$

Exponential: $N(x) = N_s \exp[-|x-x_p|/x_d]$

Gaussian: $N(x) = N_s \exp[-(x-x_p)^2/x_d^2]$

Erfc: $N(x) = N_s \text{erfc}[x/x_d]$

where N_s is the Peak Concentration,
 x_p is the Depth of Peak,
 x_d is the Profile Factor, and
 x is the distance from the specified region surface.

To remove a diffusion, set the Peak Concentration for that diffusion to zero. Also shown in the menu are the calculated junction depth and diffusion sheet resistance ($\Omega/\square$). These values are calculated based on the diffusion parameters and background doping for the region, using the current temperature and carrier mobility parameters. The calculation of these values for one diffusion ignores the presence of dopants from other diffusions. These values should therefore be interpreted with caution when there are multiple diffusions in the same region. When the diffusion is of the same type as the background doping, the net sheet resistance representing the parallel combination of the diffused region and the background is also shown.

External doping profiles are standard ASCII files with a filename of the form *.DOP. These can be created using a text editor or generated by another computer program. Each line in the file should contain three numerical values, each separated by one or more spaces. The first value on each line is a value of position within the individual region, x (cm). The second value on each line is the donor doping density, N_d , at that position (cm^{-3}), and the final value is the acceptor doping density, N_a , at that position (cm^{-3}). The position values, x , must start at 0 and increase monotonically. The values of dopant density must be nonnegative. The position values will become the initial set of nodes that separate and define the finite elements for solving the problem. PC-1D may readjust the node positions during the solution (Section 3.06). There must be at least

6 nodes defined in the file, though many more nodes will likely be needed in areas that are expected to form depletion regions or areas of rapid doping variation. The total number of positions for all three regions cannot exceed the maximum number of elements allowed, 150.

The last value of x in the external doping file does not define the thickness of the region, which is determined by a separate entry (Section 3.10). If the range of x in the file is less than the specified thickness, the last element defined in the file is expanded to include the rest of the thickness. This way, you can specify a front-surface diffusion profile and simulate the effect of varying substrate thickness without altering the doping file. You can use a separate region for the back surface if you want to have a diffusion there. If the range of x in the file exceeds the specified thickness, only part of the doping-file data is used. For regions 2 and 3, the first doping values in the file (at $x=0$) are actually assigned to a position that is offset into the region by an amount equal to the specified grading-layer thickness (Section 3.10). If this offset is greater than one-half the thickness of the first element defined in the file (equal to the second x value in the file), PC-1D will effectively reduce the grading-layer thickness so that it does not exceed one-half of the first file-defined element.

3.13 Recombination

This line indicates whether the Internal model is used for the bulk recombination lifetime or if an external disk file has been specified. Upon selecting the recombination item, you must first choose one of the currently active regions. Regions are made active by setting their thickness to a nonzero value (Section 3.10). For each region you can specify the recombination parameters for the bulk of the region and the front and back surfaces. If the internal recombination models are selected in conjunction with internal doping profiles (Section 3.12), the net low-level-injection (LLI) values for the background bulk recombination lifetime and surface recombination velocity at each surface will be displayed.

If you select bulk recombination, you are given a choice of the Internal model or an external Disk File. The internal model uses the Shockley-Read-Hall (SRH) equation with separate electron and hole lifetimes and an adjustable trap energy level within the bandgap.

$$R = \frac{pn - n_{ie}^2}{\tau_{no}(p + n_{ie}e^{-E_t/kT}) + \tau_{po}(n + n_{ie}e^{E_t/kT})} \quad (13)$$

The trap level, E_t , is specified with respect to the intrinsic level, which lies near the middle of the bandgap. If the internal doping profile is used, you can choose alternatively to specify either the net LLI background carrier lifetime or the background diffusion length. If you specify one of these values, the electron and hole lifetimes will be set equal and adjusted to provide the requested value. Because Auger and band-to-band

recombination is included in these calculations, it may not be possible to obtain the requested LLI lifetime or diffusion length even in the total absence of SRH recombination. In such cases, both lifetimes are set to their maximum value of 1 second, giving the largest possible LLI lifetime and diffusion length for the current background doping.

External bulk-lifetime files are standard ASCII files having a filename of the form *.TAU. These files contain three numerical values on each line, separated by one or more spaces. The first value on each line is a value of x , in cm. The second value on each line is the electron lifetime, τ_{no} (seconds). The final value on each line is the hole lifetime, τ_{po} (seconds). The SRH trap level is assumed to lie at the intrinsic level ($E_t = 0$). The values of x must increase monotonically, starting with $x=0$. The maximum number of lines allowed is 150. The last value of x in this file does not define the thickness of the device, which is determined by a separate entry (Section 3.10). If the range of x is less than the thickness, the final lifetime value is extended for the rest of the thickness. If the range of x exceeds the thickness, only part of the bulk-lifetime file is used. Note that the bulk-lifetime file applies to the entire device, not just one region. If a bulk-lifetime file is selected for any region, that file applies to all regions.

If you select either front or back surface recombination, you are given the choice of either a surface recombination velocity model or a saturation current density model. The surface recombination velocity model is based on the SRH equation and thus operates in a manner similar to the internal bulk-lifetime model, including the ability to directly specify the LLI net surface recombination velocity.

$$R_s = \frac{pn - n_{ie}^2}{(p + n_{ie}e^{-E_{ts}/kT})/S_n + (n + n_{ie}e^{E_{ts}/kT})/S_p} \quad (14)$$

The saturation current density model causes the recombination rate at the surface to be directly related to the excess pn product at the surface, which in turn is related to the splitting of the electron and hole quasi-Fermi potentials:

$$qR_s = J_0 (pn - n_{ie}^2)^{1/\gamma} = J_0 [e^{q(\phi_p - \phi_n)/\gamma kT} - 1], \quad (15)$$

where q is the elementary charge, $1.6021773E-19$ C. The ideality factor, γ , is included to permit the modeling of depleted-surface recombination.

When more than one region is active, the front and back surfaces refer to the individual regions, not the device as a whole. The recombination rate specified for the back surface of region 1 is added to the recombination rate for the front surface of region 2 to get the total rate of recombination at the interface between the two regions.

3.14 Surfaces

This line indicates which of three electrostatic models are to be used for the front and back surfaces of the device. The available models are Neutral, Barrier, and Charged. A Neutral surface is assumed to have no net volume charge density at the surface. A surface with a Barrier has a net charge at the surface due to band bending. You can specify the barrier height in electron-volts (eV). The bands are shifted at the surface by this amount, in a direction so that a positive value creates an energy barrier for majority carriers. This band bending results in an equilibrium volume charge density at the surface. PC-1D assumes that this equilibrium volume charge density remains unchanged in the presence of excitation. A Charged surface also results in band bending near the surface. In this case, the density of charges external to the device can be specified (charges/cm²). A positive density implies positive external charge, which produces a negatively-charged region within the device. Enter a negative surface density of charges if the external charge is negative, producing a positively-charged region within the device. Although identical equilibrium band bending can be obtained by specifying either a Barrier height or surface Charge, the two surface conditions differ when excitation is applied. While the Barrier maintains a constant surface volume charge density, the Charged surface maintains a constant surface electric field.

3.15 Circuit

This line shows which circuit elements are currently connected to the device. There are two identical sets of circuit elements; one associated with the base contact and one associated with the collector contact. Both sets contain an external Thevenin-equivalent voltage source and series resistance, and components for representing internal series resistance, shunt conductance, and a shunt diode. An additional interface shunt conductance is available, primarily for modeling tunnelling across interfaces. These circuit elements are depicted in Fig. 3. The emitter contact is always connected; but the base, collector, and interface shunt can be either Connected or Disconnected. The position of each contact within the device can be specified. You can enter a large value (i.e. 1E4 μm = 10 mm) to insure that a contact will be located at the back surface, regardless of the device thickness.

The concept of a contact in the middle of a device may seem odd at first, but this generalization is necessary to model three-terminal devices. In a three-dimensional situation, the current for the middle contact (e.g. the base of a bipolar transistor) is conducted to the active region of the device from a contact that is removed some distance perpendicular to the active device dimension. Assuming that this remote contact is in a heavily-doped area that remains in low-level injection, a voltage imposed on the contact shifts the majority-carrier quasi-Fermi potential from its equilibrium value by an amount equal to the applied voltage. This shift in quasi-Fermi level can be assumed to be nearly constant between the contact

Fig. 3 PC-1D external circuit elements.

and the active region of the device, even when the active region is in high-level injection. A small variation in the quasi-Fermi potential may be caused by resistance between the contact and the active region, but this can be modeled as an external series resistance. PC-1D assumes that all contacts can be treated this way: as injection points for majority-carrier current at which the majority-carrier quasi-Fermi potential is shifted by the applied voltage. Given this generality, it is possible to locate the contacts anywhere in the device.

The resistances R_{cc} and R_{bb} are used to represent the external power sources. They affect how much current flows in the device for a given source voltage, but the loss associated with these elements is not reflected in the device performance as revealed in plots of collector or base current versus voltage. The remaining elements, including the resistors R_c and R_b , are used to represent internal device losses. The values of these elements do affect plots of device current versus voltage.

A constant-voltage condition at a contact is achieved by setting the corresponding source series resistor to zero. A constant-current condition is maintained by setting both the source voltage and series resistance to large values, so that their ratio gives the desired current. The open-circuit voltage at either the collector or base contact can be obtained by setting the source voltage to a small (can be zero) value, and the corresponding resistance to a large value (1E6 Ω is usually plenty).

The base and collector shunt diodes and resistors can be used to incorporate three-dimensional effects found in real devices. For example, junction leakage often occurs where the junction intersects the semiconductor surface. This can usually be modeled as a combination of shunt conductance and a diode having an ideality factor of 2 or more. The shunt diode can also be used to model recombination that occurs in an area of the device not explicitly being solved. An example of this application is given in the solar cell example (Section 6.3).

The interface shunt conductance allows you to provide a path for current between two locations within the device. This is useful for modeling tunnel contacts in multiple-junction devices, where, due to heavy doping, current should flow easily across a junction. Because PC-1D does not directly include the tunneling mechanism in its equations, it is necessary to model this with a shunt conductance. The interface shunt should be used with caution, as it can adversely affect the convergence of the numerical solution (Section 4.5).

3.16 Base/Collector Voltage

These two items are used to set up the magnitude and time dependence of the voltage sources that may be connected to the collector and base contacts (Section 3.15). The choices are Zero, Fixed, Step, Sweep, and Disk File. A single source voltage can be entered for Fixed operation, while both initial and final values are required for Step and Sweep. When a transient solution begins, it uses the initial voltage values to calculate the steady-state condition. Based on the selection of a Step or Sweep function, the voltage is either stepped immediately to the final value at $t=0^+$, or else it is ramped linearly from the initial value to the final value with a slope that causes it to reach the final value on the last time step (Section 3.07).

Time-dependent voltage source values can also be entered from an external ASCII data file having a filename of the form *.VLT. Each line in the file should contain a single numerical value, a source voltage in units of volts. The first line should contain the value to be used for the steady-state solution. The second line contains the value to be used at time $t=0^+$. The remaining lines will be used by subsequent time steps. The maximum number of time steps is 50, so the maximum number of lines allowed is 52. The actual number of values used from the file is determined by the number of time steps requested (Section 3.07). If more time steps are requested than there are entries in the file, the last voltage is used for all subsequent time steps.

3.17 Light

This item is used to set the magnitude and time dependence of light incident on the front surface of the device. If you don't want to have any photogeneration of carriers, set this line to Zero. This will turn off all

of the photogeneration procedures and reduce the time required for PC-1D to set up your problem. To turn on the light source, select Fixed, Step, Sweep, or Disk File. A single light intensity can be entered for Fixed operation, while both initial and final values are required for Step and Sweep. The values entered give the total power density normally incident on the front surface of the device, measured in W/cm^2 . One standard "terrestrial sun" corresponds to an intensity of $0.1 \text{ W}/\text{cm}^2$. When a transient solution begins, it uses the initial intensity values to calculate the steady-state condition. Based on the selection of a Step or Sweep function, the light is either stepped immediately to the final value at $t=0^+$, or else it is ramped linearly from the initial value to the final value with a slope that causes it to reach the final value on the last time step (Section 3.07).

Time-dependent light intensity values can also be entered from an external ASCII data file having a filename of the form *.LGT. Each line in the file should contain a single numerical value, a light intensity in units of watts/cm^2 . The first line should contain the value to be used for the steady-state solution. The second line contains the value to be used at time $t=0^+$. The remaining lines will be used by subsequent time steps. The maximum number of time steps is 50, so the maximum number of lines allowed is 52. The actual number of values used from the file is determined by the number of time steps requested (Section 3.07). If more time steps are requested than there are entries in the file, the last voltage is used for all subsequent time steps.

3.18 Photogeneration

This item only appears if the light intensity is non-zero. The choices are Spectrum, Monochromatic, Sweep, or Disk File. The Spectrum option allows you to select one of several spectrums of incident wavelengths. These spectrums are provided as ASCII files having filenames of the form *.SPC. Three spectrums are provided with PC-1D. AM0.SPC is the solar spectrum outside of the earth's atmosphere. AM15D.SPC is the spectrum of the light that reaches the earth's surface directly from the solar disk when the sun's position in the sky causes it to travel through 1.5 standard atmospheres. AM15G.SPC also corresponds to 1.5 standard atmospheres, but includes the global diffuse blue skylight that reaches the earth's surface. These files contain two values on each line, separated by one or more spaces. The first value is a wavelength in nm, and the second is a power density in W/cm^2 (not spectral density, $\text{W}/\mu\text{m}/\text{cm}^2$). The power densities are scaled as necessary so that the total light intensity for the spectrum as a whole equals the value requested (Section 3.17). The entries must be in order of increasing wavelength, with a maximum of 100 wavelengths allowed.

The Monochromatic option allows you to specify that all of the incident power occurs at a single wavelength. The Sweep function also provides monochromatic input, but causes the wavelength to ramp linearly as a function of time from the specified initial wavelength to the final wavelength with a slope such that the final wavelength is reached on the last time step (Section 3.07).

The Spectrum, Monochromatic, and Sweep options all force PC-1D to internally calculate the photogeneration profile within the device. At each incident wavelength, after accounting for front-surface reflection (Section 3.19), the light is absorbed in the device using the absorption coefficients for each region (Section 3.11). If the device is Faceted (Section 3.09), then the photons do not travel parallel to the solution direction (x). The direction they travel, shown in Fig. 2, makes a different angle with respect to x near the front surface than near the back. Both angles are calculated using the facet angle and the index of refraction for the material of region 1, with the transition between these two angles assumed to occur abruptly when x exceeds one-sixth of the facet depth. The photons reflect from the back surface with either the same angle at which they arrived (polished) or randomly-directed (textured). If either surface is textured, light trapping occurs (Section 3.19).

The Disk File option allows you to enter the photogeneration rate as a function of position within the device directly. These files are standard ASCII files having filenames of the form *.GEN. Each line of the file contains two values, separated by one or more spaces. The first value is the distance from the front surface, in cm. The second value is the cumulative photogeneration between the front surface and this position, in units of electron-hole pairs per second. Both the position and photogeneration values must be monotonically increasing functions within the file, and both must start with a value of 0.0 on the first line. Ideally, the photogeneration values provided in the file should be scaled to correspond to 1 W/cm^2 total input area power density, because the file values will be scaled by the specified light intensity (Section 3.17) when PC-1D sets up the photogeneration profile. If the range of positions included in the file is less than the device thickness, the final photogeneration value is extended for the rest of the thickness, in effect providing no additional photogeneration beyond that position. Only part of the generation-file information is used if the device thickness is less than the position range provided in the file.

3.19 Reflectance

This line, which only appears when the light intensity is nonzero (Section 3.17), displays the values assumed for the reflectance of the front and back surfaces and whether each surface is textured. If this item is selected, you have the opportunity to modify the reflectance assumptions for either the front or back surface.

If you choose the front surface, several options for the calculation of reflectance are provided: Fixed, Coated, or Disk File. Fixed reflectance is constant for all incident wavelengths. The Coated option allows you to specify the thickness and index of refraction for up to three optical coating layers, plus a baseline reflectance that is the same for all wavelengths. Layer 1 corresponds to the first layer deposited while layer 3 is the last layer deposited. Thus light must traverse the layers in reverse numerical order as it passes from the air into the device. Table 1 gives the index of refraction for some common optical-coating materials for wavelengths in the visible spectrum. Facet angles are not taken into consideration in the calculation of reflectance for a coated surface. If your front surface is textured, you should either enter a Fixed reflectance or provide an external reflectance Disk File.

To minimize reflectance, the layers should be stacked in order of decreasing index of refraction with the highest-index material adjacent to the device. The thickness of each layer should be one-quarter wavelength within the material. A layer of Si_3N_4 would need to be about 70 nm thick to minimize reflectance at a wavelength of 600 nm. Note that some coating materials may absorb light, especially in the ultra-violet. PC-1D does not take this effect into account.

Table 1
Commonly Encountered Indices of Refraction
Room Temperature, Solid Form, at 589.3 nm

SiO_2	silicon dioxide	1.46
Si_3N_4	silicon nitride	2.05
ZnS	zinc sulfide	2.36
MgF_2	magnesium fluoride	1.38
Al_2O_3	aluminum oxide	1.76
TiO_2	titanium dioxide	2.62
Ta_2O_5	tantalum pentoxide	2.20
	organic films	≈ 1.4

External reflection files are standard ASCII files with a filename of the form *.REF. Each line in the file should contain two numerical values, separated by one or more spaces. The first value on each line is a value of wavelength, in nm. The second value on each line is the reflectance. The maximum number of lines allowed is 100. All values of the reflectance must be between 0 and 1, and the values of wavelength must increase monotonically. If the range of wavelengths provided is less than is required for the solution of a problem, the reflectance for the first (or last) wavelength is used for all smaller (or larger) wavelengths.

If you choose the back surface, your only option is to enter a fixed reflectance, independent of wavelength. This is the internal reflectance for light that reaches the back surface as a result of entering the front of the device and traveling its entire thickness. Some of these reflected photons return to the front surface, where they may be either internally

reflected or transmitted out of the device.

You can specify whether each surface is polished or randomly texturized. If both surfaces are polished, the internal reflectance of the front surface is equal to its external reflectance to incident light. If either the front or back surface is texturized, light-trapping occurs. In this case you must specify a front-surface internal reflectance for randomly-directed photons. The ideal-case effective internal reflectance for randomly-oriented light is $1 - (1/n^2)$, where n is the index of refraction for region 1.

This section describes the operation of PC-1D after the problem has been fully specified using the interactive input mode described in Section 3. Section 4.1, "Solution Phases," explains the steps taken by PC-1D to obtain a valid solution of the semiconductor equations, and how you can interact with this process. Section 4.2, "Examination Mode," describes how to use the interactive examination mode to explore in detail the results of the computation. Section 4.3, "User-Defined Graphs," provides additional details about this feature of the interactive examination mode. Section 4.4, "Batch Mode," describes how to solve a whole series of problems unattended. Section 4.5, "Convergence Failure," gives a few tips on how to proceed when your problem fails to converge to a solution.

As soon as a problem has been fully specified using the interactive input mode (Section 3), it is ready for solution. Before starting the solution, it is a good idea to check the device thickness, doping profiles, lifetimes, mobilities, optical parameters, and functions of time using the Preview Parameters option (Section 3.02) to invoke the interactive examination mode (Section 4.2). There are so many input parameters that it is very easy to set up a different problem than what you had intended.

4.1 Solution Phases

The solution involves three distinct phases: equilibrium (EQ), steady state (SS), and transient (TR). Equilibrium and steady state each require one solution of the semiconductor equations, while the transient phase requires one solution of the equations for each time step (including $t=0^+$). In addition, an EQ solution is performed after every renode operation (Section 5.4). Table 2 illustrates which phases are performed, in what order, based on the requested solution mode (Section 3.07) and the phase that was most recently completed. If reinitialization is requested, as it normally is (Section 3.05), then the last phase completed is reset to None and the solution simply progresses through the phases (starting with EQ) until the solution mode requested is completed.

Table 2
Solution Phases

<u>Solution Mode</u>	<u>Last Phase Completed</u>	<u>Phases Attempted</u>
Equilibrium	None, EQ, SS, or TR	EQ
Steady State	None	EQ, SS
Steady State	EQ, SS, or TR	SS
Transient	None	EQ, SS, TR
Transient	EQ	SS, TR
Transient	SS	SS, TR
Transient	TR	TR (extend number of steps)

The program will complete the solution of the phases indicated in Table 2 without interruption unless one of the following occurs:

1. The user presses the Enter key (Examine) during the solution of any phase. An asterisk will appear in the center of the convergence indicator bar. When the current solution converges, the program enters the interactive examination mode (Section 4.2) to permit detailed evaluation of the solution variables. On exiting from examination mode, the solution will continue where it left off.
2. The user presses the Esc key (Input) during the solution of any phase. The program will return directly to the interactive input mode. This is used to interrupt the solution if the user discovers an erroneous input parameter, or if an adjustment of the numerical parameters (Section 3.06) is needed to speed convergence.
3. Convergence Failure. If the solution diverges to a relative error exceeding $1E20$ or if the number of iterations taken exceeds the maximum specified by the input parameters (Section 3.06), then the solution is interrupted as if the Esc key had been pressed.

During each attempt to solve the semiconductor equations, a convergence indicator bar is drawn at the bottom center of the screen. Shown to either side of this indicator are reminders of the action of Enter and Esc, and text indicating the last successfully completed phase or time step. The convergence indicator displays the maximum change calculated for any of the three solution potentials (Section 5.1) resulting from the iteration just completed. The maximum change implemented may be somewhat less due to clamping (Section 3.06). The convergence scale is logarithmic, with the centerline indicating a maximum change in potential (normalized by the thermal voltage, kT/q) of 1.0. The left edge corresponds to successful convergence as defined by the specified error limit (Section 3.06). The right edge is the inverse of the error limit. Thus for an normalized error limit of $1E-6$, the right edge indicates a maximum calculated error of $1E6$ and the left edge indicates $1E-6$.

As each solution of the equations is obtained, four graphs are plotted on the display, as shown in Fig. 4. These graphs are selectable by the user from the set of predefined graphs (Section 3.04). The default graphs are carrier concentration, electrostatic potential, current density, and the band diagram.

4.2 Examination Mode

An interactive graphical Examination Mode is provided to allow the user to study the problem solutions in detail, dump the solution data to a disk file, or generate hard-copy output using a pen plotter. This mode is entered when one of the following occurs:

Fig. 4 Default four-graph display.

1. The program successfully completes all phases required for the specified solution mode (Section 4.1), and the user presses Enter in response to the prompt at the bottom of the screen.
2. The user presses the Enter key during the problem solution, the current solution successfully converges (Section 4.1), and the user presses Enter in response to the prompt at the bottom of the screen.
3. The user selects either the Preview Parameters or Examine Results options from the PC-1D Program Control menu (Section 3.02), and presses Enter in response to the prompt at the bottom of the screen.

In the case of Preview Parameters, the solution variables are initialized to reflect the values that will be used as the starting point for the next problem solution. The following functions of position are set up and available for previewing: doping, mobility, cross-sectional area, dielectric constant, recombination lifetime, intrinsic concentration, and bandgap. The spatial photogeneration rate is not available to preview; it is calculated at the start of each solution phase. The following functions of time are also set: collector voltage source, base voltage source, and light intensity. If the illumination is swept between two wavelengths, then the wavelength, front-surface reflection, index of refraction, and absorption coefficient in each region are set. The index of refraction and absorption coefficients can only be previewed for region 1. For steady-state solutions these time functions provide only a single point at $t=0^-$. The escaped-light reflectance is not available to preview; it is calculated in conjunction with the photogeneration profile. Some of the functions can only be previewed as a User-Defined Graph (Section 4.3). None of the

solution variables (carrier concentrations, electric field, currents, voltages, etc.) have any validity within the Preview Parameters option.

In all situations other than Preview Parameters, Examination Mode allows the user to explore the values of spatial variables resulting from the last four solutions completed, and the values of temporal functions for all time steps that have been calculated. Many of these functions are only accessible as a User-Defined Graph (Section 4.3), including all spatial information for solutions prior to the one most recently completed. If the last solution attempted was interrupted by either a Convergence Failure (Section 4.5) or by pressing the Esc key (Section 4.1), the values of spatial variables for that solution will not represent a valid solution of the semiconductor equations.

Upon entering Examination Mode, you will be presented with a menu of predefined graphs. Or, you can select the User-Defined graph, described in Section 4.3. When you select one of these graphs, a full-screen plot of that information will be displayed. The graph may include from one to four separate sets of data, each plotted as a curve having a different linetype (solid, dashed, dotted, and dot-dashed). A small menu will appear in the upper-right corner of the screen with the cursor positioned to the first entry. This menu offers you three options: Enter Trace Mode, Dump Data to Disk, or Hard-Copy Plot.

If you Enter Trace Mode, a cursor will appear at the left side of the graph, positioned on the solid-line curve. At the bottom of the screen you will see either two or three data values. Only two data values are needed if the abscissa (x-axis) of the graph is either "Distance from the Front" or "Time." When two values are displayed, the one on the left is the abscissa value at the current cursor location and the value on the right is the ordinate (y-axis) value at the current cursor location. Other graph options require the display of three data values. The value on the left is the value of the parameter. The parameter will either be distance from the front (x), or time (t). The center value is the abscissa value and the value on the right is the ordinate value for the current cursor position. For these parametric plots, both the abscissa and ordinate are functions of position or time, so that for each value of the parameter (spatial or temporal), there is a single value for the abscissa and a single value for the ordinate. When there are multiple curves, each curve provides its own ordinate value for each value of the parameter. The most common parametric plots are temporal: current versus voltage, and response versus wavelength. But interesting spatial options are available, like mobility versus doping, and lifetime versus carrier concentration. Fig. 5 shows an example of a spatial parametric graph.

Fig. 5 Examination mode parametric graph.

In trace mode, the cursor is moved around the graph using either the mouse or the cursor keys. Movement to the right increases the parameter, movement to the left decreases it. For graphs where the abscissa is the parameter (position or time), this causes the cursor to move right or left as expected. For parametric plots, however, the direction the cursor takes may not correspond to the direction of the mouse or cursor keys. The Home key can be used to move the cursor to the lowest parameter value, and the End key to the highest parameter value.

The cursor can be switched to a different curve by using the up and down cursor keys, or a vertical movement of the mouse. An upward movement will switch the cursor to the next-higher curve, and a downward movement will switch the cursor to the next-lower curve. Once on the new curve, left and right movements can be used to explore the data values. Note that when the cursor changes curves, the label for the ordinate data value (on the right side at the bottom on the screen) changes to correspond to the new curve. This data label can be used to determine which curve is which when there are several related curves on one plot. Refer to Tables 3 and 4 in Section 4.3 for a list of the available functions and their corresponding data labels. It is also helpful to recognize that, when there are two curves, one for each carrier, the electron curve is solid and the hole curve has holes in it.

It is often desirable to examine a subset of the data plotted by enlarging one portion. To accomplish this, move the cursor to one side of the region to be enlarged (the cursor can be on any curve if there is more than one), press the Enter key, move the cursor to the opposite side of the region to be enlarged, and press Enter again. The data between these two points,

inclusive, will be replotted on a new set of axes scaled to fit the reduced data set. At this point you will again receive the small menu offering the chance to Enter Trace Mode, Dump Data to Disk, or Hard-Copy Plot. You can reenter trace mode, examine the data as before, and even use the Enter key to enlarge a subset of the enlargement! In fact, you can use this recursive feature to enlarge portions of enlargements up to four consecutive times. Each time you exit a graph by pressing the Esc key, you will be returned to the previous graph from which the enlargement was made. Hit the Esc key while in the original graph to return to the menu of predefined graphs.

The Dump Data to Disk option will create a file that contains the data currently on the screen, with a header of auxiliary information like labels and titles so that the program PC1DPLOT (Section 7) can be used to recreate the graph at some later time. If an enlarged portion of the graph is currently on the screen, only that portion of the data will be stored in the file. You will be presented with a list of plot-data files (*.PDF) in the current disk directory. You can select one of these, in which case the file selected will be overwritten with the new data. Or, you can select New File and enter a new filename. PC-1D will check to make sure that the filename you enter does not already exist. The suffix PDF is assumed and should not be typed when entering the new filename.

The Hard-Copy Plot option allows the user to generate a hard-copy plot of the data currently on the screen using a pen plotter, if one is connected and its device driver has been properly installed (Section 2). Because these paper plots are generated using a separate device driver, the resolution of the paper plot is much better than the display's resolution. This may cause minor differences in the way the axes are scaled and labeled.

4.3 User-Defined Graphs

In addition to the predefined plots for examining the results of the solution, you can design your own plots by separately selecting variables for each axis. One abscissa and up to four ordinate data sets can be plotted on a single set of axes. These data sets may be different solution variables, or the same variable from different points in time, as long as all of the data sets have the same parametric dependency (i.e. all spatial or all temporal). Title and axes labels will be generated automatically if there is only one curve (an abscissa and one ordinate data set). If there is more than one curve, then the user must supply a title (if desired), and the ordinate axis will be labeled with units only if all of the ordinate data sets have the same units. The scaling of the abscissa and ordinate axes can each be either linear or logarithmic. Fig. 6 shows the menu that is used to set up a user-defined graph.

Fig. 6 User-defined graph menu.

The parametric dependency of the plot is determined by the abscissa variable. Only ordinate variables with matching parametric dependency are permitted. After selecting an abscissa, only those variables with the correct parametric dependency will be included in the list of available functions from which to choose ordinate variables. To remove a data set from the user-defined graph, first select that data set using the Enter key. The Available Functions menu will appear. Rather than selecting a data set, press Esc. That data set will be removed. If that data set was the abscissa, then all four ordinate data sets are also removed, leaving a clean slate.

Tables 3 and 4 list all of the spatial and temporal functions that are available, including the corresponding data labels and comments. Spatial information is available for the current solution and the three previous solutions, making it possible to graph the evolution of a spatial variable over four time steps.

4.4 Batch Mode

Unattended batch-mode operation is supported, making it possible to perform optimization studies with PC-1D that may require hundreds of solutions and several days of uninterrupted computation. First, each problem is set up using the interactive input mode (Section 3). Rather than solve each problem as it is set up, the parameters which fully describe the problem can be stored as a parameter file (*.PRM) on the disk (Section 3.03). This process can be repeated for a large number of related problems, storing each set of parameters in a separate file. A batch file (*.BAT) is then created which calls PC-1D repeatedly, each time specifying a different

Table 3
Available Spatial Functions
"#" in unit labels is replaced by an appropriate SI prefix.

Function	Label	Units	Comments
Distance from Front	x	#m	
Electron Concentration	n	cm-3	
Hole Concentration	p	cm-3	
Electric Potential	Psi	#V	
Electric Field	E	#V/cm	
Charge Density	Rho	#C/cm3	
Electron Current Density	J _n	#A/cm2	
Hole Current Density	J _p	#A/cm2	
Total Current Density	J _T	#A/cm2	
Donor Doping Density	N _D	cm-3	
Acceptor Doping Density	N _A	cm-3	
Band Gap	E _g	#eV	Includes bandgap narrowing
Intrinsic Concentration	n _{ie}	cm-3	Includes bandgap narrowing
Minority Carrier Lifetime	Tau	#sec	Includes Auger, band-band
Recombination Rate	R	cm-3/s	Bulk recombination only
Cumulative Recombination	R _c	pairs/s	Includes surface recomb.
Photogeneration Rate	G	cm-3/s	
Cumulative Photogeneration	G _c	pairs/s	Spatial integral of G(x)
Electron Velocity	v _n	cm/s	J _n /(qn)
Hole Velocity	v _p	cm/s	J _p /(qp)
Electron Mobility	u _n	cm2/Vs	Includes velocity saturation
Hole Mobility	u _p	cm2/Vs	Includes velocity saturation
Excess Electron Conc.	Δn	cm-3	n - n _{eq}
Excess Hole Concentration	Δp	cm-3	p - p _{eq}
Excess Potential	ΔPsi	#V	ψ - ψ _{eq}
Excess Charge Density	ΔRho	#C/cm3	ρ - ρ _{eq}
Electron Quasi-Fermi Energy	E _{Fn}	#eV	Zero ref. is E _{Fmaj} at x=0
Hole Quasi-Fermi Energy	E _{Fp}	#eV	Zero ref. is E _{Fmaj} at x=0
Conduction Band Edge	E _c	#eV	Zero ref. is E _{Fmaj} at x=0
Valence Band Edge	E _v	#eV	Zero ref. is E _{Fmaj} at x=0
Vacuum Energy	E _{vac}	#eV	Zero ref. is E _{Fmaj} at x=0
Resistivity	Rho	#ohm*cm	
Conductivity	Sigma	#S/cm	
Cumulative Conductivity	Sigma _c	#S	Spatial integral of Sigma(x)
Carrier Diffusion Length	L _D	um	Includes velocity saturation
Electron Current	I _n	#A	
Hole Current	I _p	#A	
Total Current	I	#A	
Electron Drift Current	I _{dn}	#A	qμ _n n(∇ψ+∇V _n)A
Hole Drift Current	I _{dp}	#A	qμ _p p(∇ψ+∇V _p)A
Electron Diffusion Current	I _{Dn}	#A	qD _n ∇nA
Hole Diffusion Current	I _{Dp}	#A	-qD _p ∇pA
Cross-Sectional Area	A	cm2	
Dielectric Constant	k		Relative DC permittivity

Table 4
Available Temporal Functions
"#" in unit labels is replaced by an appropriate SI prefix.

Function	Label	Units	Comments
Time	t	#sec	t=0 ⁻ is steady-state
Free-space Wavelength	W	nm	
Spectral Response	SR	#A/W	$I_b/(P_L \cdot A)$
External Quantum Efficiency	QE _{ex}	%	$I_b \cdot hc/(P_L \cdot A \cdot q \cdot \lambda)$
Internal Quantum Efficiency	QE _{in}	%	$I_b \cdot hc/(P_L \cdot A \cdot q \cdot \lambda \cdot (1 - R_f - R_e))$
Front Reflectance	R _f	%	% of incident beam energy
Escape Reflectance	R _e	%	% of incident beam energy
Incident Light Intensity	P _L	#W/cm ²	
Incident Light Energy	E _{ph}	#eV	Energy per photon, hc/λ
Collector Current	I _c	#A	Includes D _c , G _c current
Collector Voltage	V _c	#V	Includes drop across R _c
Base Current	I _b	#A	Includes D _b , G _b current
Base Voltage	V _b	#V	Includes drop across R _b
Collector Source Voltage	V _{cc}	#V	
Base Source Voltage	V _{bb}	#V	
Base Input Power	P _{b-1}	#W	$V_b \cdot I_b$
Inverse QE	QE ⁻¹		$1/QE_{in}$
Absorption Length	L _{abs}	#m	$1/\alpha$, For region 1
Absorption Coefficient	alpha	cm ⁻¹	For region 1
Collector Input Power	P _c	#W	$V_c \cdot I_c$
Index of Refraction	n _r		For region 1

parameter file for input and a different set of files to accept the output results. These output files are in the same format (*.PDF) that is generated by the Dump Data to Disk option of the Examination Mode (Section 4.2). The data that is placed in these files is defined by the selection of the output graphs (Section 3.04). The first output file listed receives the graph data selected for the upper-left screen, the second file receives the graph data for the upper-right screen, the third receives the graph data for the lower-left screen, and the fourth receives the graph data for the lower-right screen. Only predefined graphs are available for batch-mode output. If fewer than four output filenames are listed, then fewer graphs are saved. These output files can be examined, submitted to a user-written program for evaluation, or plotted using PC1DPLOT (Section 7).

Suppose, for example, that three problems have been set up using the interactive input mode and stored with the filenames PROB1.PRM, PROB2.PRM, and PROB3.PRM. We are only interested in the graphs of collector current versus voltage and the electron and hole concentrations versus position for the last time step, so these two graphs were specified in each problem for the upper-left and upper-right screens, respectively. The batch file PROBS.BAT to run all three problems consecutively might look like this:

PC-1D PROB1 PROB1IV PROB1NP
PC-1D PROB2 PROB2IV PROB2NP
PC-1D PROB3 PROB3IV PROB3NP

The choice of output filenames is arbitrary. If no suffix is given, the default suffix of PDF will be appended. This batch file would be executed simply by typing PROBS in response to the DOS prompt. When it finishes, we should have six output files full of the desired data. If one or more of the output files is missing, it means that particular solution failed to converge (Section 4.5). Note that only spatial information from the last time step in a transient solution can be saved. If you want to record spatial information at other time steps, you need to create additional PRM files that stop after fewer time steps, and add corresponding lines to the batch file. If you want to stop a batch-mode solution while it is in progress, press Ctrl-Break.

4.5 Convergence Failure

The equilibrium solution of the semiconductor equations is remarkably robust and will almost always converge. But the extreme nonlinearity of the fully-coupled semiconductor device equations makes convergence to a nonequilibrium solution difficult. Several measures have been taken within PC-1D to assist convergence, but despite this you will eventually generate problems that won't converge. This section offers some advice on how to avoid convergence failure, and how to deal with it when it does occur.

When the normalized error exceeds $1E20$ or the maximum number of iterations permitted is exceeded, a Convergence Failure is reported. The solution process is terminated and, if you are not in batch mode, you are returned to the PC-1D Program Control menu. You will have to modify the problem parameters to improve the convergence properties. Check first for the following items, which are known to interfere with convergence.

1. Setting Reinitialize to No after changing a significant material parameter such as the thickness, doping, or temperature.
2. Locating an electrical contact in a lightly doped or depleted region.
3. Applying a current source in reverse bias.
4. Applying two different voltage sources in close proximity.
5. Connecting the interface shunt between two dopant regions that are both isolated by a junction from another electrical contact.
6. Applying a large forward voltage with no current-limiting resistor.
7. Negligible recombination. Turning off all recombination in the device (surface and bulk) gives the minority carriers no place to go.

Most often, the device description is perfectly legitimate and the convergence problem is due to the solution trying to reach a final state that is too far removed from the initial state. The abrupt application of voltage, current, or light can all cause trouble. The problem is worsened in high-bandgap materials, where the minority carrier density may change by twenty orders of magnitude when moderate excitation is applied. In these situations you can try adjusting the normalized potential clamp (Section 3.06), although a smaller clamp value will not always improve the situation. Adjusting the element size factor (Section 3.06) affects the accuracy of the solution but has only a minor impact on the convergence properties. By far the most effective way of improving convergence is to impose the excitation gradually. For example, to solve a silicon junction forward biased to 0.8 volts, you may need to perform interim solutions at 0.6 and 0.7 volts first. Similarly, it may be necessary to increase the light intensity in steps, say 1 mW/cm^2 , 10 mW/cm^2 , then 100 mW/cm^2 . After each interim solution be sure to set Reinitialize to No so that the interim solution will be used as the starting point for the next solution.

Convergence of transient solutions can be particularly troublesome, especially for short time steps. The worst case is between the steady-state condition at $t=0^-$ and the start of the transient at $t=0^+$. To represent this brief interval, PC-1D takes a finite time step equal to 1% of the specified time step size (Section 3.07). If there is a step change in any of the excitations, it occurs during this step. As a general rule, convergence is better for sudden removal of excitation than for sudden application. If you have only swept excitations there is no change in the excitation parameters between steady-state ($t=0^-$) and the start of the transient ($t=0^+$), so you would expect to see no change in the solution. But because time-dependent terms are added to the semiconductor equations for the transient, you may encounter a convergence problem at $t=0^+$ even if the solution converged for steady state. Use a larger step size if possible when this situation arises. If your intent is to use the transient capability to solve for a progression of essentially steady-state solutions, then don't hesitate to use step sizes in the ksec range if necessary.

This section describes the isolated-element approach used within PC-1D to obtain rapid, accurate solutions of the semiconductor transport equations. These two goals are far from compatible. Increasing the number of elements provides greater accuracy, but the time required increases proportionately. The emphasis in this program has been on reliable convergence for a broad range of doping profiles and boundary conditions.

The five semiconductor equations (Section 5.1) can be readily reduced to three equations in three unknowns. There are several options for selecting the solution variables, none appears to have a significant advantage. For PC-1D/3 we have chosen the electrostatic potential, ψ , and the electron and hole quasi-Fermi potentials, ϕ_n and ϕ_p . Three expressions are required at each interface between elements. Poisson's equation minimizes the total electrostatic energy in the adjoining elements, and both electron and hole currents must be continuous across each interface.

5.1 Fundamental Equations

The numerical method of PC-1D uses the two-carrier semiclassical semiconductor transport equations. These equations are derived from the Boltzmann transport equation with the following assumptions: the two carriers flow independently (no carrier-carrier scattering), both carrier populations remain in thermal equilibrium with the surrounding crystal lattice (no hot carriers), the mobility of carriers is isotropic, and the structure of the energy levels available to electrons is not significantly affected by excitation (rigid bands). The effect of magnetic fields is also neglected and the device temperature is assumed to be uniform. The resulting equations for electron and hole current density (J_n, J_p) are commonly written in terms of the electron and hole density (n, p), the electron and hole mobility (μ_n, μ_p), and the electron and hole quasi-Fermi energies (E_{Fn}, E_{Fp}).

$$\begin{aligned}\bar{J}_n &= \mu_n n \nabla E_{Fn} \\ \bar{J}_p &= \mu_p p \nabla E_{Fp}\end{aligned}\tag{16}$$

The quasi-Fermi energies relative to their corresponding energy band define which of the energy states available to electrons are filled by carriers from each population, thereby determining the local concentration of both carriers. For nondegenerate material, the Boltzmann exponential approximation to the Fermi distribution function can be utilized:

$$\begin{aligned}n &= N_c e^{-(E_c - E_{Fn})/kT} \\ p &= N_v e^{-(E_{Fp} - E_v)/kT}\end{aligned}\tag{17}$$

Here N_c and N_v are the effective density of states in the conduction and valence bands. E_c and E_v are the energy levels of the two band edges. The effects of degenerate doping can be included in this Boltzmann-like form by modifying E_c and E_v . Degeneracy due to high injection levels is not included in this treatment.

Define the quasi-Fermi potentials (ϕ_n, ϕ_p) as the separation of the quasi-Fermi energies relative to the equilibrium Fermi energy at some selected reference location. For PC-1D, this location is taken as $x=0$. The position of the two band edges in Eq. (17) relative to this reference energy depends on the electrostatic potential and any spatial variation in the band structure. To describe variations in the band structure, two band-edge potentials are introduced, V_n and V_p . These potentials represent the shift (toward the center of the bandgap) of the conduction and valence band edges with respect to the conditions at the reference location, after accounting for electrostatic effects. These shifts may be caused by spatial variations in material composition or bandgap narrowing due to heavy doping. Specifically, for bandgap E_g , electron affinity χ , effective conduction (valence) band density of states N_c (N_v), and conduction (valence) bandgap narrowing ΔE_{gc} (ΔE_{gv}),

$$\begin{aligned} V_n &= (kT/q) \ln(N_c/N_{cr}) + (\chi - \chi_r)/q + \Delta E_{gc}, \text{ and} \\ V_p &= (kT/q) \ln(N_v/N_{vr}) - (E_g - E_{gr})/q - (\chi - \chi_r)/q + \Delta E_{gv}, \end{aligned} \quad (18)$$

where q is the magnitude of the elementary electronic charge. The material (conduction and valence band) properties at the reference location can serve to establish a convenient zero reference for the electrostatic potential. Let zero electrostatic potential correspond to neutral intrinsic conditions in this material. Introducing these definitions into Eq. (17) produces a new pair of Boltzmann-like expressions for n and p ,

$$\begin{aligned} n &= n_{ir} e^{q(\psi + V_n - \phi_n)/kT} \\ p &= n_{ir} e^{q(-\psi + V_p + \phi_p)/kT} \end{aligned} \quad (19)$$

Here, n_{ir} is the intrinsic carrier concentration at the reference location.

Combining Eqs. (16) and (19) it is possible to eliminate n and p , expressing the two unknown current densities in terms of the two unknown quasi-Fermi potentials and the unknown electrostatic potential. All other parameters can be expressed in terms of these unknown quantities for a given device structure, providing two equations for five unknowns.

Two additional equations can be generated by accounting for the fate of all electrons and holes that enter a given volume of space. These continuity equations account for flow into the volume, creation of carriers within the volume (generation), and annihilation of carriers within the volume (recombination).

$$\begin{aligned}
 \frac{\partial n}{\partial t} &= \frac{\nabla \cdot \bar{J}_n}{q} + G_L - U_n \\
 \frac{\partial p}{\partial t} &= - \frac{\nabla \cdot \bar{J}_p}{q} + G_L - U_p
 \end{aligned}
 \tag{20}$$

Here G_L is the volume generation rate of electron-hole pairs due to the absorption of light (photogeneration), and U_n , U_p are the net volume recombination rates for electrons and holes.

The photogeneration model in PC-1D includes a number of sophisticated features (Section 3.18), including light trapping. It does not consider the effect of heavy doping on optical absorption, the generation of multiple carrier pairs by energetic photons, two-photon absorption processes that use a trap level as an intermediate step, or absorption of photons created during radiative recombination. PC-1D also does not include generation of carriers due to impact ionization, and thus will not predict avalanche breakdown.

A consequence of the two-carrier model is that electrons and holes are always created and annihilated in pairs, so that $U_n = U_p = U$. This approximation is valid only for high-quality material in which the concentration of occupied trap states is small compared with the carrier populations. The net recombination (U) can, in turn, be expressed in terms of the three unknown potentials (ψ , ϕ_n , ϕ_p). The recombination model includes trap-assisted, Auger, and band-to-band processes (Section 3.13). Other factors must be neglected for reasons of computational efficiency, including changes in the trap cross-sectional area or energy level as a function of occupation state, and the effect of traps distributed throughout the bandgap.

The addition of Eq. (20) makes it possible to eliminate J_n and J_p as unknowns, leaving two equations in the three unknown potentials. Poisson's equation of electrostatics provides the third and final independent equation.

$$\nabla \cdot (\epsilon \nabla \psi) = - \rho \tag{21}$$

where ρ is the volume charge density and ϵ is the permittivity. Contributions to the charge density include the electron and hole densities and the concentration of ionized donor and acceptor impurities. The two-carrier model neglects the charge associated with trap states located within the bandgap. PC-1D assumes that all of the available dopants are ionized. This is an adequate assumption at room temperature and above, especially since the true dopant concentration is rarely known with great accuracy. At lower temperatures, the user must adjust the doping profiles to represent only the ionized dopant density.

As with any set of differential equations, boundary conditions represent an essential component in defining a unique solution. For the semiconductor transport equations, three boundary conditions are required at each external surface. A variety of boundary conditions can be used. Most analytical models assume that (1) the majority-carrier concentration is fixed at its equilibrium value, (2) the minority-carrier concentration is related to the minority-carrier current density by a surface recombination velocity, and (3) the surface electrostatic potential is shifted from its equilibrium value by an amount equal to the externally-applied voltage. These conditions are appropriate when all surfaces are heavily doped and electrically contacted, either directly or indirectly.

A more general set of boundary conditions is used in PC-1D to accommodate lightly doped and insulated surfaces. The three boundary conditions at each surface are based on (1) injected current density, (2) surface recombination, and (3) surface electrostatic potential. The injected total current density is zero at insulated surfaces, and is controlled by the external circuit at contacts. PC-1D provides two surface recombination models: a Shockley-Read-Hall recombination velocity model that permits separate electron and hole velocities through a trap level located anywhere in the bandgap, and a saturation current density model that gives a recombination rate proportional to the excess pn product at the surface. Three boundary conditions are available for the surface electrostatic potential: neutral surface charge density, the surface charge density remains at its equilibrium value, or the surface electric field remains at its equilibrium value. Constant surface charge density is appropriate at metal-semiconductor contacts, including Schottky contacts, whereas constant surface electric field is more appropriate at insulated surfaces.

At electrical contacts, PC-1D allows the amount of injected current to be determined, in part, by the device being simulated. A Thevenin-equivalent circuit is attached to each electrical contact, consisting of a voltage source, series resistance, and external shunt elements. By varying the source resistance from a small value to a large value, the entire spectrum of boundary conditions ranging from fixed-voltage to fixed-current can be simulated.

A complication in applying the injected-current boundary condition arises when the contact itself is not a part of the solution region. It is often the case that a significant distance may lie between the contact and the solution region. Contacts to practical devices are made to heavily-doped material, so one can generally assume that the injected current will enter the solution region as majority-carrier flow spread throughout the region that matches the dopant type at the contact. To simplify matters, PC-1D concentrates this majority-carrier current into a single injection point. The resulting electrostatic potential shift at the point of injection presumably causes a similar shift in the electrostatic potential near the contact. There will be a substantial difference, however, if the point of current injection in the solution region is in high-level injection. The

difference in injection levels between the solution region and the contact produces an electrostatic potential difference across the intervening material, even when no current is flowing. For this reason, it is better to use the majority-carrier quasi-Fermi potential to determine the voltage at contacts. In the heavily-doped material adjacent to the contact, the majority-carrier quasi-Fermi potential must shift in unison with any shift in the electrostatic potential in order to maintain a fixed volume charge density at the surface. Thus, the majority-carrier quasi-Fermi level is a valid indicator of the contact voltage, with the advantage that it does not differ significantly between the solution region and the contact. What difference might exist will be proportional to current and can therefore be modeled using the external series resistance.

5.2 Discretization

The first step in the numerical solution of the semiconductor equations described in Section 5.1 is the discretization of the device, both in time and space. PC-1D is a quasi-one-dimensional model, so the device is spatially divided into a finite set of M elements along a solution dimension "x", which approximately follows the direction of current flow in the device and is not necessarily a straight line. The dividing point between any two elements is called a node. The nodes are numbered from zero at the left boundary to M at the right. Subscripts of j or k are used to denote values at the j^{th} or k^{th} nodes, respectively. The cross-sectional area for current flow is a function of x ; with a value A_k assigned to each node.

Poisson's equation is satisfied within this discrete space of elements by minimizing the total electrostatic energy of the system, W_e .

$$W_e = \int_0^w \frac{\epsilon}{2} \left[\frac{\partial \psi}{\partial x} \right]^2 dx - \int_0^w \rho \psi dx, \quad (22)$$

where w is the width (thickness) of the device. A minimization of this functional with respect to ψ , for a given charge distribution, ρ , gives the electrostatic potential. This minimization is achieved in a discrete sense by approximating ψ , ρ , and ϵ . Both ψ and ρ are given a linear spatial variation within each element. The permittivity, ϵ , is assumed to be uniform within each element. Setting the differential energy to zero yields the following discrete equation at each internal node k :

Poisson's Equation:

$$\epsilon^- \frac{A_{k-1} + A_k}{2} \frac{\Delta\psi_{k-1}}{\Delta x_{k-1}} - \epsilon^+ \frac{A_k + A_{k+1}}{2} \frac{\Delta\psi_k}{\Delta x_k} = \frac{\rho_{k-1} \Delta x_{k-1} A_{k-1} + 2\rho_k (\Delta x_{k-1} + \Delta x_k) A_k + \rho_{k+1} \Delta x_k A_{k+1}}{6}, \quad (23)$$

where ϵ^+ and ϵ^- refer to the permittivity of the elements on the right and left sides of the node, respectively. In Eq. (23), a delta notation is used to refer to the difference in a value across an element; that is, $\Delta\psi_k = \psi_{k+1} - \psi_k$, and $\Delta\psi_{k-1} = \psi_k - \psi_{k-1}$.

A second discrete equation at each node is based on the transport equation for electrons. Start with a quasi-one-dimensional version of Eq. (16), which gives the electron current, I_n , in the direction of the solution dimension.

$$I_n(x) \bar{a}_x = -q \mu_n(x, E) n(x) A(x) \nabla \phi_n(x). \quad (24)$$

The mobility in Eq. (24) is electric-field dependent. The model used for this dependency is

$$\mu(x, E) = \mu(x, 0) / \sqrt{1 + [\mu(x, 0)E/v_{\text{sat}}(x)]^2}, \quad (25)$$

where $E = -\nabla\psi$ is the electric field intensity and v_{sat} is the high-field saturation velocity, which may be different for electrons and holes. The denominator of Eq. (25) is assumed to be uniform within each element.

Define a modified electric potential, $\xi_n(x)$.

$$\xi_n(x) = \psi(x) + V_n(x) + \frac{kT}{q} \ln \frac{\mu_n(x, E) A(x)}{\mu_r A_r}, \quad (26)$$

where μ_r and A_r are arbitrary fixed reference values for mobility and area. Substitute the potentials of Eq. (19) for n in Eq. (24) and rewrite the expression in terms of the modified potential of Eq. (26).

$$I_n(x) \bar{a}_x = -q \mu_r n_{ir} A_r e^{q(\xi_n(x) - \phi_n(x))/kT} \nabla \phi_n(x). \quad (27)$$

Eq. (27) can be rearranged so that the quasi-Fermi potential is the only position-dependent value on the RHS of the equation.

$$I_n(x) e^{-q\xi_n(x)/kT} \bar{a}_x = qD_r n_{ir} A_r \nabla e^{-q\phi_n(x)/kT}, \quad (28)$$

where $D_r = \mu_r \cdot kT/q$ defines the reference diffusivity. A contour, or path, integral can be applied to both sides of this equation. The contour is taken across a single element, following the solution dimension x from the node at x_k to x_{k+1} .

$$\int_c I_n(x) e^{-q\xi_n(x)/kT} \frac{-q\xi_n(x)}{a_x} \cdot \overline{dl} = qD_r n_i r A_r \int_c \nabla e^{-q\phi_n(x)/kT} \cdot \overline{dl}. \quad (29)$$

$$\int_{x_k}^{x_{k+1}} I_n(x) e^{-q\xi_n(x)/kT} dx = qD_r n_i r A_r [e^{-q\phi_n(x_k)/kT} - e^{-q\phi_n(x_{k+1})/kT}]. \quad (30)$$

The value of this approach is now apparent. To complete the integration on the LHS of Eq. (30), it is only necessary to make assumptions regarding the spatial variation of two well-behaved variables: I_n and ξ_n . No assumption regarding the spatial variation of ϕ_n was necessary. Since the carrier concentration, n , depends on both ξ_n and ϕ_n , this is equivalent to saying that no constraints have been imposed on the spatial variation of n within the element. If we assume that both ξ_n and I_n vary linearly within each element, the integral can be performed to produce a direct expression for the electron current at the left edge of the element (node x_k). Letting β be the inverse thermal voltage, q/kT , the following discrete equations result.

Electron Current Transport:

$$I_n(x_k) = -\Delta I_n Y(\beta\Delta\xi_n) + \frac{qD_r n_i r A_r}{\Delta x} [e^{\beta(\xi_n(x_{k+1}) - \phi_n(x_{k+1}))} Z(\beta\Delta\xi_n) - e^{\beta(\xi_n(x_k) - \phi_n(x_k))} Z(-\beta\Delta\xi_n)], \quad (31)$$

$$\text{where } Z(u) = \frac{u}{e^u - 1} \text{ and } Y(u) = \frac{1 - Z(u)}{u}.$$

In Eq. (31), the net electron current leaving the element, ΔI_n , is equal to the total rate of recombination within the element, minus the total rate of generation, plus the rate that electrons are accumulating as a function of time. In PC-1D, photogeneration within each element is based on assuming a separate exponential absorption characteristic for each wavelength, the recombination rate is assumed to vary linearly with position across each element, and carrier accumulation is calculated assuming that the excess electron concentration is a linear function of position. The time dependence of the accumulation is calculated as a backward difference,

assuming an exponential function of time unless the excess has changed sign since the last time step, in which case a linear time dependence is used.

A similar derivation gives an expression for hole current, providing the third discrete equation needed at each node.

Hole Current Transport:

$$I_p(x_k) = -\Delta I_p Y(-\beta\Delta\xi_p) - \frac{qD_r n_{ir} A_r}{\Delta x} [e^{-\beta(\xi_p(x_{k+1}) - \phi_p(x_{k+1}))} Z(-\beta\Delta\xi_p) - e^{-\beta(\xi_p(x_k) - \phi_p(x_k))} Z(\beta\Delta\xi_p)]. \quad (32)$$

5.3 System Solving

The electron and hole currents at node k can each be calculated independently for the two elements that adjoin at that node. Each carrier current should be continuous at the node except for interface recombination and majority-carrier current that may be injected at that node from an external source. Using the discrete expressions for carrier current developed in Section 5.2, two equations at each internal node k result.

Hole Current Transport:

$$\begin{aligned} & - \frac{D_r}{\Delta x_{k-1}} [e^{-\beta(\xi_p - \phi_p)_k} Z(-\beta\Delta\xi_p)_{k-1} - e^{-\beta(\xi_p - \phi_p)_{k-1}} Z(\beta\Delta\xi_p)_{k-1}] \\ & + \frac{(\Delta I_p)_{k-1}}{n_{ir} A_r} Y(\beta\Delta\xi_p)_{k-1} = \\ & - \frac{D_r}{\Delta x_k} [e^{-\beta(\xi_p - \phi_p)_{k+1}} Z(-\beta\Delta\xi_p)_k - e^{-\beta(\xi_p - \phi_p)_k} Z(\beta\Delta\xi_p)_k] \\ & - \frac{(\Delta I_p)_k}{n_{ir} A_r} Y(-\beta\Delta\xi_p)_k + \frac{R_s}{n_{ir}} - \frac{I_{maj}}{qn_{ir} A_r}. \end{aligned} \quad (33)$$

Electron Current Transport:

$$\begin{aligned}
 & \frac{D_r}{\Delta x_{k-1}} [e^{\beta(\xi_n - \phi_n)_k} Z(\beta \Delta \xi_n)_{k-1} - e^{\beta(\xi_n - \phi_n)_{k-1}} Z(-\beta \Delta \xi_n)_{k-1}] \\
 & + \frac{(\Delta I_n)_{k-1}}{n_{ir} A_r} Y(-\beta \Delta \xi_n)_{k-1} = \\
 & \frac{D_r}{\Delta x_k} [e^{\beta(\xi_n - \phi_n)_{k+1}} Z(\beta \Delta \xi_n)_k - e^{\beta(\xi_n - \phi_n)_k} Z(-\beta \Delta \xi_n)_k] \\
 & - \frac{(\Delta I_n)_k}{n_{ir} A_r} Y(\beta \Delta \xi_n)_k - \frac{R_s}{n_{ir}} - \frac{I_{maj}}{qn_{ir} A_r} .
 \end{aligned} \tag{34}$$

Eqs. (33) and (34), combined with Eq. (23), provide three simultaneous nonlinear equations at each internal node. Added to these are three equations at each boundary that are based on (1) injected current density, (2) surface recombination, and (3) surface electrostatic potential. The first two of these boundary conditions are implemented using Eqs. (33) and (34), with the expression for current that extends beyond the boundaries of the device set to zero (i.e. set LHS to zero at $x=0$). The electrostatic-potential boundary condition either equates the volume charge density at the surface to its equilibrium value (0 for a neutral surface), or equates the electric field at the surface to its equilibrium value. For M elements, there are a total of $3(M+1)$ equations for $3(M+1)$ unknowns.

These nonlinear equations are solved using Newton's method, which iteratively solves a sequence of linearized approximations of the equations. A Jacobian matrix is generated using the derivatives of each equation in terms of ϕ_n , ψ , and ϕ_p , each normalized by kT/q . The derivatives are calculated from analytical expressions for all terms except the surface and bulk recombination, which are calculated numerically based on a 1% change in the excess carrier densities. These expressions are such that the equations for node k involve only values from nodes $k-1$, k , and $k+1$. The resulting matrix has a numerical bandwidth of 11. The three equations are ordered so that the matrix will be diagonally dominant, eliminating the need for time-consuming pivoting. The electron current continuity equation is diagonal with ϕ_n , the hole current continuity equation is diagonal with ϕ_p , and the electrostatic equation is diagonal with ψ . A solution vector is constructed using the residual error from the $3(M+1)$ system equations. The entries for a single node are shown in Eq. (35).

$$\begin{vmatrix} \frac{\partial f_1}{\partial \phi_n^-} & \frac{\partial f_1}{\partial \psi^-} & \frac{\partial f_1}{\partial \phi_p^-} & \frac{\partial f_1}{\partial \phi_n^0} & \frac{\partial f_1}{\partial \psi^0} & \frac{\partial f_1}{\partial \phi_p^0} & \frac{\partial f_1}{\partial \phi_n^+} & \frac{\partial f_1}{\partial \psi^+} & \frac{\partial f_1}{\partial \phi_p^+} \\ \frac{\partial f_2}{\partial \phi_n^-} & \frac{\partial f_2}{\partial \psi^-} & \frac{\partial f_2}{\partial \phi_p^-} & \frac{\partial f_2}{\partial \phi_n^0} & \frac{\partial f_2}{\partial \psi^0} & \frac{\partial f_2}{\partial \phi_p^0} & \frac{\partial f_2}{\partial \phi_n^+} & \frac{\partial f_2}{\partial \psi^+} & \frac{\partial f_2}{\partial \phi_p^+} \\ \frac{\partial f_3}{\partial \phi_n^-} & \frac{\partial f_3}{\partial \psi^-} & \frac{\partial f_3}{\partial \phi_p^-} & \frac{\partial f_3}{\partial \phi_n^0} & \frac{\partial f_3}{\partial \psi^0} & \frac{\partial f_3}{\partial \phi_p^0} & \frac{\partial f_3}{\partial \phi_n^+} & \frac{\partial f_3}{\partial \psi^+} & \frac{\partial f_3}{\partial \phi_p^+} \end{vmatrix} \begin{vmatrix} \Delta \phi_n^0 \\ \Delta \psi^0 \\ \Delta \phi_p^0 \end{vmatrix} = - \begin{vmatrix} f_1 \\ f_2 \\ f_3 \end{vmatrix}, \quad (35)$$

where superscripts $^-$, 0 , and $^+$ correspond to values at nodes $k-1$, k , and $k+1$. The diagonal matrix components are shown in bold characters.

The Jacobian matrix is inverted using LU decomposition. Multiplying the factored Jacobian matrix times the solution vector generates a set of values for $\Delta \phi_n$, $\Delta \psi$, and $\Delta \phi_p$ at each node. The electrostatic potential at each node is then updated using the following function, which puts a clamp on the amount the potential is allowed to change during a single iteration. The clamp value is supplied by the user. The default is 5 (kT/q).

$$\Delta \psi' = \Delta \psi / (1 + |\Delta \psi / \text{clamp}|). \quad (36)$$

The changes in the two quasi-Fermi potentials at each node are also limited, but in a different manner. In addition to clamping the change in each potential according to Eq. (36), the separation between each quasi-Fermi potential and the electrostatic potential is also monitored and prevented from changing by more than the clamp value during each iteration.

$$\Delta \phi' = \Delta \psi' + (\Delta \phi - \Delta \psi') / (1 + |(\Delta \phi - \Delta \psi') / \text{clamp}|). \quad (37)$$

At the end of each Newton iteration the maximum normalized update (prior to clamping) is determined. The magnitude of this correction is displayed on the screen as a way of monitoring the convergence of the solution. The solution continues until this maximum normalized correction error falls below the user-supplied error limit, typically 1.0E-6.

For equilibrium solutions, the knowledge that $I_n = I_p = G_L = R = 0$ at all x allows the semiconductor equations to be reduced to a single equation in the single unknown, ψ . Thus for equilibrium there are only $M+1$ equations in $M+1$ unknowns, and the numerical bandwidth is reduced to 3. The equations are still quite nonlinear, and the same iterative Newton method is applied. At the end of each iteration ψ is updated, with the clamping function of Eq. (36) applied. The iterations continue until the maximum change in $\beta \psi$ is less than the user-supplied normalized error limit (Section 3.06).

5.4 Node Placement

The flexibility of PC-1D is greatly enhanced by its ability to automatically place nodes where they are needed most. This "renoding" takes place at various times as requested by the solution parameters (Section 3.06). After renoding, the error will usually increase for a few iterations, then converge rapidly toward zero.

The node placement algorithm used in PC-1D does not guarantee a specific number of elements. Rather, the user supplies an "Element Size Factor" which directly reflects the intrinsic accuracy of the solution. For the same value of this factor, some problems may require many more elements than others. If more than 150 elements are necessary to achieve the specified level of accuracy, PC-1D will automatically increase the element size factor until the number of elements drops below 150. The default value for this factor (α) is 1.0. For more accuracy, consider using a smaller value of α , like 0.5 or 0.25. Values greater than 1.0 provide faster operation, but the solution accuracy will deteriorate.

The node placement algorithm starts with an initial set of elements and proceeds to scan them in order of increasing x . For each element, PC-1D chooses one of three options: (1) keep the element as it is, (2) subdivide the element into up to 30 smaller elements, or (3) concatenate the element to the element on its left. The determination of which option to take depends on several criteria. These criteria are based on the fundamental assumptions in the numerical method and on producing visually pleasing graphs.

An element will be concatenated to the one on its left unless any of the following tests are true for the two elements combined.

- (1) The average volume charge density causes the electrostatic potential to deviate from linearity by an excessive amount:

$$\left\{ \frac{\rho \Delta x^2}{2 \epsilon} \right\}^{1/2} > \alpha kT/q. \quad (38)$$

- (2) The majority-carrier doping concentration changes too rapidly:

$$|\Delta \ln(N_{\text{maj}})| > \alpha. \quad (39)$$

- (3) The width of the element is too large for its proximity to the surface (This is to insure an adequate number of small elements near the front surface in case rapidly-absorbed light is subsequently applied):

$$\Delta x > x. \quad (40)$$

- (4) The element consumes too large a fraction of the width of the region:

$$\Delta x > w_{\text{reg}}/20. \quad (41)$$

- (5) The node that would be eliminated by concatenation is the interface between regions or an injection point for the emitter, base, collector, or interface shunt contacts.

The element will be subdivided if either criterion (1) or (4) is true. The number of subelements created will be equal to the ratio of the LHS to the RHS of the equation, with a maximum of 30 subelements permitted.

After a new set of nodes has been selected, the values of the various solution variables at these nodes is obtained by interpolating values from the previous set of nodes. For N_D , N_A , τ_{no} , and τ_{po} , this interpolation assumes that the logarithm of the parameters is linear between nodes; while for ψ , ψ_{eq} , ϕ_n , ϕ_p , A , V_n , V_p , ξ_n , ξ_p , v_{sat} , E_t , ϵ , and G it assumes that the values themselves are linear.

This section presents a series of four different devices and leads you through the use of PC-1D to examine their operation. The devices are presented in order of increasing solution complexity, from the simple analysis of an equilibrium inversion layer to the generation of heterojunction bipolar transistor characteristic curves. The calculated results are compared to analytical theory in several instances where the theory should give an accurate prediction. This should give you the confidence to apply PC-1D to cases where the simple theory is not applicable. Working your way through these examples is also an excellent way to introduce yourself to PC-1D's capabilities.

If you have not rebooted your system since you installed your CONFIG.SYS file, do so now, following the instructions in Section 2. Invoke the program by typing PC-1D in response to the DOS prompt. You can later exit the program by pressing Esc from anywhere in the Program Control menu. Once you get into the interactive-input mode, you can start with any of the device examples which follow; it is not necessary to go through all of them in order. If you need to quit the program part-way through an example problem, be sure to either Revise or Update the parameters (Section 3.03) before you exit to DOS. That way, PC-1D will start up next time where you left off.

6.1 MOS Inversion Layer

This simulation examines the inversion layer which is formed when an applied electric field repels majority carriers and attracts minority carriers to the surface of a semiconductor. This is the situation which occurs beneath the gate in an MOS transistor. For this simulation we consider the simple case of a uniformly-doped n-type substrate which has been inverted to create a p-type channel at its surface.

Move the cursor to the Simulation Parameters line and press Enter. Press enter again to indicate that you want to load a new set of parameters. Select MOS.PRM from the list of filenames. If MOS is not on the list, it means that you didn't copy the example problems properly and should repeat the installation process of Section 2.

You should feel free to move the cursor around the Program Control menu. You can press Enter on the Material, Doping, or Surfaces lines to examine some of the relevant problem parameters, but don't change any of them yet. In particular, note that the Surfaces line has been set to "Charged" for the front surface. An external surface charge density of $-1\text{E}11$ has already been entered for the front surface, corresponding to $1\text{E}11$ electrons per cm^2 . Return to the Program Control menu (press Esc as many times as required if you are in a submenu). To begin the solution, move the cursor to the top of the menu (Proceed with Solution) and press Enter.

As this is just an equilibrium solution, you won't have to wait very long before the solution is complete. Press Enter as directed by the note at the bottom of the screen to obtain the menu of graphs. Select the Volume Charge Density graph. It should appear as shown in Fig. 7.

Fig. 7 MOS inversion layer charge density.

Notice the distinct surface inversion region and the extended depletion region which adjoins it. Note the smooth edge of this region, in contrast to the abrupt edge which is assumed in the simplified depletion-region approximation. The depletion approximation predicts a depletion region width under strong inversion of

$$w_D = \sqrt{\frac{4\epsilon kT \ln(N_D/n_i)}{q^2 N_D}} \quad (42)$$

Using the parameters for this problem ($N_D = 1.0\text{E}15 \text{ cm}^{-3}$, $\epsilon = 11.9 \epsilon_0$, $n_i = 1.01\text{E}10 \text{ cm}^{-3}$, and $kT/q = 0.02585 \text{ V}$), this expression predicts a width of $0.88 \text{ } \mu\text{m}$, which does correspond approximately to the middle of the gradual transition at the edge of the calculated space-charge region.

Press Enter (or the left mouse button) to enter the trace mode. You can use the left and right cursor keys (or mouse) to move the cursor and read the actual values from the graph. You might want to verify that the charge density in the depletion region is equal to $q \cdot N_D$. When you are done, press Esc (or the second mouse button) to return to the graph menu. You can examine any of the available graphs if you wish. Press Esc while in the graph menu to return to the Program Control menu. Note that the completed solution state is now Equilibrium. For practice, return to graphics mode

by selecting the Preview / Examine option from the Program Control menu, then specifying that you want to Examine Results. Press Esc to return to the Program Control menu when you are finished.

6.2 Forward-Biased Diode

This simulation introduces a number of additional complexities. A nonuniform doping profile is used to create an n^+p asymmetrical diode. This introduces position-dependent mobilities and bandgap narrowing. The peak electric field in the junction region is examined and compared to theory. A forward bias voltage is applied, causing the minority carrier densities to increase greatly and produce a net device current. Finally, the device's turn-off transient in a specific circuit is examined, illustrating how current can flow backwards through a diode until it regains its reverse blocking capability. The time constant associated with this transient is related to the junction capacitance and compared with analytical theory.

Move the cursor in the Program Control menu to the Simulation Parameters line and press Enter. Press Enter again to load a new set of device parameters. Select DIODE.PRM from the list of filenames. After the device parameters have been loaded from the disk, move the cursor to the Doping line and press Enter. Select region 1 (the only region in this device), then select the Internal Model. This device has a heavily-doped n-type diffusion into a p-type background. Select the front diffusion and confirm that the emitter has a complementary-error-function (Erfc) profile with a surface doping of $1E18 \text{ cm}^{-3}$ and a junction depth of $1.0 \text{ } \mu\text{m}$. If you were to use a four-point probe to measure the sheet resistance on the front surface, you would measure $536 \text{ } \Omega/\square$. Esc back to the Doping menu and select the background doping. It should be $1.0E16 \text{ cm}^{-3}$, corresponding to a resistivity of $1.4 \text{ } \Omega\text{cm}$. Esc back to the Program Control menu.

It is a good idea to check to make certain that the problem has been set up properly before beginning the solution. Move the cursor to the Preview / Examine line and press Enter. Select the Preview Parameters option. After the four graphs have been drawn, press Enter as directed at the bottom of the screen to obtain the menu of available graphs. Select Doping Profiles from this list. You should get a full-screen plot of the doping profiles similar to Fig. 8. The device should be $4 \text{ } \mu\text{m}$ thick, with a junction at $1 \text{ } \mu\text{m}$. Press Enter to select the trace mode. Using the cursor (or mouse) to move around the graph, you should be able to determine that the doping density gradient, "a", near the junction is $6.9E20 \text{ cm}^{-4}$. To proceed with the solution, press Esc as many times as necessary to return to the Program Control menu, make sure the cursor is on the top line, and press Enter.

When the equilibrium solution is complete, observe the four plots shown. The upper left plot of carrier concentrations clearly shows the depletion layer separating the emitter and base of the device. There is no current in this device because we have not yet applied any voltage.

Fig. 8 Diode doping profile.

Press Enter to obtain the menu of graphs. Select the Electric Field and obtain a graph similar to Fig. 9. We can compare the peak electric field calculated by the program with an analytical approximation. Based on the theory of the linearly-graded junction, we would expect to have a space-charge layer width of

$$W = \left[\frac{12\epsilon V_g}{q a} \right]^{1/3}, \quad (43)$$

where the gradient voltage, V_g (the electrostatic potential across the space-charge region), is defined in Eq. (44).

$$V_g = \frac{2}{3} \frac{kT}{q} \ln \left[\frac{a^2 \epsilon kT}{8q^2 n_i^3} \right] = 635 \text{ mV}. \quad (44)$$

Inserting this value of V_g into Eq. (43) gives a space-charge layer width of $0.417 \mu\text{m}$. The peak electric field should be approximately

$$E_{\text{max}} = \frac{qaW^2}{8\epsilon} = 23 \text{ kV/cm}. \quad (45)$$

Select trace mode and confirm that the calculated peak value of the electric field is about 23 kV/cm. Press Esc to return to the Program Control menu when you are finished.

Fig. 9 Equilibrium diode electric field.

Note that on the Circuit line, the base voltage source, V_{bb} , has been connected. From the Base Voltage line you can see that the base voltage has been set to 0.6 volts. You can investigate the other parameters, also. From the Recombination line you can determine that the surface recombination velocity is $1E3$ cm/s at the front and $1E7$ cm/s at the back surface. $1E7$ cm/s corresponds to an essentially infinite recombination rate at the back surface, such that the velocity is limited only by the thermal velocity of the carriers. The low-level-injection (LLI) recombination lifetime in the substrate (the base region) is set for $10 \mu s$, which is long enough to ensure that virtually all recombination will take place at the rear surface rather than in the bulk.

To solve for the application of forward bias, move the cursor to the Solution Mode line, press Enter, and select the Steady State option. Then move the cursor to the top of the Program Control menu and press Enter. The program will quickly recompute the equilibrium solution, then begin the steady-state solution. You will note that each steady-state iteration takes somewhat longer than for equilibrium. When the solution is complete you should note from the upper-left graph that the forward-bias carrier concentrations are many orders of magnitude higher than they were in equilibrium, while the majority carrier density is essentially unchanged. Press Enter to get the graph menu, then select the Carrier Concentrations option to obtain a graph similar to Fig. 10. Select the trace mode for closer examination.

Fig. 10 Forward-biased diode carrier concentrations.

Note that at this voltage the device is still in low-level injection. Near the junction, simple analytical theory suggests that we should expect to have

$$np = n_i^2 e^{qV/kT} = 1.23E30 \text{ cm}^{-6}. \quad (46)$$

At $x = 1.0134 \text{ } \mu\text{m}$, the cursor readout shows $n = 1.16E15 \text{ cm}^{-3}$ and $p = 1.05E15 \text{ cm}^{-3}$, giving a calculated $np = 1.22E30 \text{ cm}^{-6}$, essentially identical to the analytical result.

Assuming that this device can be treated as an asymmetrical short-base diode ($W = 3 \text{ } \mu\text{m}$) in low-level injection with no space-charge layer recombination yields an analytical expression for the current density of

$$J = \frac{qD_n n_i^2}{W N_A} e^{qV/kT} = 2.06 \text{ A/cm}^2, \quad (47)$$

where $D_n = 31.5 \text{ cm}^2/\text{s}$ was obtained from the graph of mobility and multiplying by kT/q . For the given device area of 0.01 cm^2 , this corresponds to a device current of 20.6 mA. The PC-1D computer calculation gives a device current of 21.2 mA (See the base current vs. voltage graph). After examining the various graphs for this steady-state solution, return to the Program Control menu.

It's time to try a transient solution. The transient to be investigated corresponds to the circuit diagram shown in Fig. 11. When the voltage source switches from 5 Volts to 0 Volts, a simple approach would suggest that the current should immediately stop flowing. In fact, it will reverse direction and flow backwards through the diode for a short period of time. To set up this circuit in PC-1D we need to add the series resistor and set up the time dependence of the base voltage source. Select the Circuit line of the Program Control menu. Select the Base Contact, then select Connected. On the first line of the Base Circuit menu, enter a source series resistance of $100\ \Omega$. Esc back to the Program Control menu. Select the Base Voltage line, then request a Step voltage change. Set the initial voltage to 5 volts, and the final voltage to 0 volts. Esc back to the Program Control menu. In order to solve the transient problem, you must select the Solution Mode line and request a Transient. You must then specify the size and number of time steps to be taken. The steps should be 10 ns ($1\text{E-}8$ sec), and 10 steps should be adequate. The number of solutions per step can stay at 1. Esc back to the Program Control menu.

Fig. 11 Diode turn-off transient circuit.

Because you have changed only the excitation of the device and not the device itself, you can save some time by moving the cursor to the Reinitialize line and pressing Enter. By setting Reinitialize to No, PC-1D will use the previous steady-state solution as the starting point for converging to the next solution.

This is a good time to specify which graphs should be displayed during the solution of the problem. Select the Output Graphs line, select Specified, and confirm that the upper-right graph is Base Current vs. Time and the lower-right graph is Base Voltage vs. Time. Esc back to the Program Control menu and proceed with the solution by selecting the top item in the menu.

The solution will proceed to step through time. As soon as the voltage is turned off, the base current will jump negative then proceed to decay toward zero. The base voltage does not jump suddenly at $t=0^+$, but instead decays toward zero from its initial steady-state value. A reasonable analytical approach to this situation is to treat the time-dependent diode as a junction capacitance in parallel with a steady-state diode. If the capacitance was constant, this would produce an exponential decay of both voltage and current with a time constant of $\tau = RC$, where the important resistance here is the $100\ \Omega$ series resistor. But we expect the junction capacitance of the diode to change depending on the voltage across it, because the width of the space-charge layer is affected. Late in the transient, however, the capacitance can be calculated using the equilibrium width of the space-charge layer from Eq. (43) to be 253 pF. This gives a time constant of 25 ns, with a slower decay (larger time constant) expected early in the transient.

When the solution is finished, you can use the User-Defined graph option (Section 4.3) to create a graph of log base current versus time as shown in Fig. 12. From the slope of the curve produced by this graph, you can determine that the range of time constants predicted by PC-1D is 63 ns to 25 ns, in good agreement with the simple theory for this case.

6.3 Solar Cell Efficiency

This simulation subjects a state-of-the-art silicon solar cell to one "sun" of solar illumination, and also calculates its response as a function of wavelength. Although a solar cell is basically a diode with light shining on it, achieving high efficiency generally requires that the cell have a rather thick base region and a thin emitter region. An additional feature of most high-efficiency cells is the presence of a diffused-dopant region at the back surface to help contain the photogenerated carriers. This device will be used to introduce the use of PC-1D's transient capabilities to scan a device's I-V curve and calculate spectral response. The efficiency of the cell will be determined from its illuminated I-V curve.

Move the cursor to the Simulation Parameters line, press Enter, then press Enter again to Load new parameters. Select CELL.PRM from the list of filenames. The parameters for an n^+pp^+ silicon solar cell will be read from the disk. You can investigate each of these parameters as they are discussed in the following paragraphs. This cell is $300\ \mu\text{m}$ thick and has a total area of $4\ \text{cm}^2$. This simulation will be performed at 25°C rather than 300 K, in keeping with standard practice for photovoltaic devices. The

Fig. 12 Diode turn-off transient.

doping profile for this device includes a $0.2 \Omega\text{cm}$ p-type base with a $100 \Omega/\square$ n-type emitter diffusion and an aluminum-alloy back-surface region. The recombination parameters for this cell are a bulk diffusion length of $200 \mu\text{m}$, front-surface recombination velocity of $1\text{E}4 \text{ cm/s}$, and a back-surface recombination velocity of $1\text{E}7 \text{ cm/s}$. The high back-surface value is because the entire back surface is assumed to be covered with metal. The front surface of the cell has narrow metal gridlines covering about 3% of its surface. The $1\text{E}4 \text{ cm/s}$ front-surface recombination velocity is a value commonly assumed to approximate the combined effect of the low-recombination passivated surface and the high-recombination area in contact with the metal gridlines. A more accurate way of treating the nonuniform front surface will be presented later in this example.

The Circuit for this simulation includes the base voltage source, a source resistor, a device resistor, and a shunt diode, as shown in Fig 13. The source series resistor is selected to limit the current from the source, and has no effect on the device performance. Its value is selected as the anticipated open-circuit voltage divided by the anticipated short-circuit current. This selection of the source series resistance produces a smooth I-V curve by concentrating the data points in the region near the knee of the curve, as illustrated by the series of load lines in Fig. 13. A source resistance of 5Ω was selected for this simulation. The device series resistor is added to model extrinsic series resistance in the cell, including sheet resistance loss associated with current flowing laterally in the emitter layer to reach a gridline, the resistance of the gridlines themselves, and contact resistance between the emitter and the gridlines. The resistance of the silicon base is included in the solution region and thus should not be added to the extrinsic series resistance. A series

resistance of $100\text{ m}\Omega$ is used for this device. The shunt diode is added to model recombination that occurs where the junction intersects the periphery of the cell. This diode has an ideality factor of 2 and a saturation current of 40 nA .

Fig. 13 Solar cell I-V circuit with curve-smoothing source resistor.

The base source voltage is set to sweep from -0.8 to $+0.8$ volts, thereby tracing out the entire illuminated I-V curve as shown in Fig. 13. The illumination is set for 0.1 W/cm^2 , which is one standard terrestrial "sun". The AM15G spectrum is used, corresponding to the global light spectrum incident on the earth's surface when the sun's angle from overhead is 48° , at which point the sunlight must pass through 1.5 times as much atmosphere as if it were directly overhead.

The cell's reflectance is based on a 3% loss due to the gridlines, plus the wavelength-dependent reflection from the polished silicon surface that follows the application of a double-layer antireflection coating of zinc sulfide and magnesium fluoride over a 6-nm surface oxide. The alloyed rear surface is textured, with an internal reflectance of about 85%. Any scattered light that returns to the front surface will see a net reflectance of 92%.

The Solution Mode for this simulation is Transient. The time step size has been set to 1 second. This is slow enough to insure that the I-V curve represents a series of essentially steady-state data points. 16 time steps are the minimum needed to generate a smooth I-V curve.

Begin the solution by moving the cursor to the top of the Program Control menu and pressing Enter. In addition to the usual parameter loading, this problem requires optical calculations. Some of these are done immediately, while the actual calculation of the photogeneration profile will be done just before the steady-state solution. The calculation of the photogeneration profile may take a significant period of time; patience is required. If you want to pause during the course of the solution to look at the variables in detail, just press Enter. An asterisk will appear in the convergence bar, and when convergence is achieved you will have the opportunity to examine the various graphs. When you Esc from the graph menu, the solution will resume where it left off.

When the I-V scan is completed, press Enter to get the Graph Menu and select Base Current vs. Voltage. The graph will appear as in Fig. 14. The cursor keys (or mouse) correspond to moving forward (right) or backward (left) in time, which allows you to review the entire curve. Also, note that the $t=0^-$ and $t=0^+$ solutions are two separate points which happen to fall directly on top of each other in this case.

Fig. 14 Solar cell I-V curve.

Start by finding the open-circuit voltage. Although you can visually estimate its value by looking for the point where the curve crosses the voltage axis, you can do better if you magnify the region near the cross-over. Move the cursor to the data point just above the axis and press Enter. Then move the cursor one position so that it is just below the axis and press Enter again. The cross-over region will be magnified, making it easy to see that the open-circuit voltage is 660 mV. Exit from the magnified region by pressing the Esc key. Similarly, you can determine that the short-circuit current where the curve crosses the current axis is

141 mA.

Esc back to the Graph Menu and select I(b) and P(b) vs. V(b). No axis label is shown for the ordinate because the two curves (current and power) have different units (amps and watts). Move the cursor to the dotted power curve using the up/down cursor keys (or mouse). Step the cursor through the entire curve, noting the device power shown at the bottom-right corner of the screen. This is the power into the device, so negative values correspond to power produced by the device. The data point with the greatest power output occurs near the knee of the curve, and gives an output of 74.8 mW. The light power incident on the cell is 400 mW ($0.1 \text{ W/cm}^2 \cdot 4 \text{ cm}^2$), so the cell's efficiency at this point on the curve is 18.7%. It is likely, however, that the point on the curve with the greatest power output (the maximum-power point) lies in-between the data points that were calculated. By assuming that the power output is approximately parabolic with respect to the source voltage near the maximum-power point, you can infer the maximum power from the three calculated points near the maximum-power point. If P_a and P_b are the values to either side of the center value, P_o , then

$$P_{\max} = P_o + \frac{(P_a - P_b)^2 / 8}{2P_o - P_a - P_b} \quad (48)$$

For the 16-point scan of this device, the correction is negligible. The fill factor is a commonly used figure of merit for solar cells, defined as

$$FF = \frac{P_{\max}}{V_{oc} I_{sc}} = 0.804 \quad (49)$$

Esc back to the Program Control menu. To calculate the spectral response (and quantum efficiency) of this cell, the excitation sources must be redefined. First, move the cursor to the Circuit line, press Enter, then select the Base Contact. Leave the base contact connected, but remove the curve-smoothing source resistor by setting R_{bb} to 0. The rest of the parameters represent the cell itself and should not be changed. Esc back to the Program Control menu. Move the cursor to the Base Voltage line, press Enter, and select the Zero option. By removing the source resistor and turning off the voltage source, you have placed a short-circuit across the cell. This is the appropriate connection for measuring spectral response. Esc back to the Program Control menu. Move the cursor to the Photogeneration line, press Enter, and select the Sweep option. The initial wavelength should be 350 nm, and the final wavelength should be 1150 nm. Esc back to the Program Control menu.

Wavelength scanning allows you to determine the spectral response of the cell (amperes of output per watt of monochromatic light) and the quantum efficiency (electrons out per photon of light). The quantum efficiency definition can be based either on all photons incident (external QE) or only on those photons that are not reflected from the cell (internal QE). The difference between internal and external QE represents the reflectance of the front surface plus any light that reflects off the back surface and then escapes out the front surface. The Quantum Efficiency graph displays all four QE components (QE_{ext} , QE_{int} , R_f , and R_e). The spectral response graph is available only as a user-defined graph.

The wavelength scan is a Transient solution. 16 steps is adequate, and has the advantage that it produces wavelength steps that are an even 50 nm. Having set up the transient wavelength scan, you can use the Preview Parameters option with user-defined graphs to examine the reflectance, index of refraction, or absorption coefficient versus wavelength.

Move the cursor to the top of the Program Control menu and press Enter to begin the solution. When the solution is complete, press Enter to get the Graph Menu and select the Quantum Efficiency graph. It should appear as shown in Fig. 15. You could generate this graph with photon energy rather than wavelength for the abscissa by using the user-defined graph option. Quantum efficiency is commonly used to infer whether a cell's short-circuit current is being reduced by recombination in the emitter (poor internal blue response) or the base (poor internal red response).

Fig. 15 Solar cell quantum efficiency versus wavelength.

It was mentioned early in this example that the assumption of an average front-surface recombination velocity could be avoided. A more accurate approach is to divide the cell into two volumes. Let the portion of the cell that is covered by gridlines constitute one volume, and the illuminated remainder be the other volume. These two volumes operate as if they were two nearly-independent devices in parallel. The volume under the gridlines is shadowed from light, so it can be represented as a normal (dark) diode, while the illuminated portion looks like an unshaded solar cell. First, PC-1D can be used to determine the parameters of the dark diode. The resulting saturation current and ideality factor can then be introduced as external circuit elements in parallel with the remaining unshaded solar cell.

To evaluate the dark diode, set the front-surface recombination velocity to $1\text{E}7$ cm/s and the light intensity to Zero. Disconnect the base contact, and connect the collector contact. This removes the circuit elements which were set previously to represent the cell, but which would interfere with the dark-diode simulation. Set the collector voltage to a fixed value of 0.65 volts (approximately the open-circuit voltage), and set the device area to a uniform 0.12 cm^2 (3% of 4 cm^2). Select a Steady-State solution mode, and proceed with the solution. When the solution is complete, determine from the Collector Current vs. Time graph that the device current is 8.35 mA. The saturation current for the dark diode in the vicinity of this operating point is given by

$$I_0 = I_{\text{dark}} e^{-qV/\gamma kT}, \quad (50)$$

where $kT/q = 25.69$ mV at 25°C . For this device the recombination occurs almost entirely in areas of low-level injection, so the ideality factor, γ , can be assumed to be 1.0. The resulting I_0 for this dark diode is 86 fA.

Esc back to the Program Control menu. To model the remaining solar cell, set the area to a uniform 3.88 cm^2 , the front-surface recombination velocity to $1\text{E}3$ cm/s (an appropriate value for this level of surface doping), and the light intensity to a fixed 0.1 W/cm^2 . Reconnect the base contact and restore the $5\ \Omega$ curve-smoothing resistor (R_{bb}), but leave the collector contact connected. Insert the 86E-15 A value for I_0 into the collector-contact circuit, and set the collector source resistor to a large value (i.e. $1\text{E}9\ \Omega$). This has the effect of connecting the dark diode across the cell while effectively isolating the collector voltage source. Set the solution mode to Transient (16 steps) and set the Base Voltage to sweep from -0.8 to +0.8 volts. Set the collector voltage to Zero. Return Photogeneration to the AM15G spectrum. The front-surface reflectance should remain Coated, but the baseline reflectance should be set to 0, since the gridline area has been excluded from the device. Proceed to solve this improved cell model, and determine for yourself that the cell's efficiency is 19.2%.

6.4 Heterojunction Transistor

This simulation explores the gain characteristics of a heterojunction NPN bipolar junction transistor. There are two phases. In the first phase, the simulation holds V_{ce} constant while V_{be} is scanned from 0.7 to 1.4 volts. By simultaneously measuring the collector and base current, a Gummel plot showing the transistor gain under a variety of base drive conditions can be graphed. In the second phase, the base current is set to 1 mA and V_{ce} is scanned from 0 to 1 volt. This produces the common-emitter output characteristic for the transistor for one base current. The complete set of common-emitter curves could be generated by repeating this simulation for several different base currents.

Move the cursor to the Simulation Parameters line, press Enter, then press Enter again to load new parameters. Select the file HBJT.PRM from the list of filenames. This is a small-area device, with two different materials. The emitter is $Al_{0.3}Ga_{0.7}As$, with GaAs for the rest of the device, giving a total thickness of 6 μm . Check the Doping parameters to see that the GaAs region has a uniformly-doped p-type base layer on an n-type collector background. The base contact location corresponds to a point midway between the emitter-base junction and the base-collector junction. For this simulation we assume that there is an interface recombination velocity between materials, as measured from the GaAs side, of $1E4$ cm/s. The front and rear surfaces have $1E7$ cm/s. The bulk recombination lifetimes are 10 ns in the GaAs and 0.2 ns in the AlGaAs. An internal base series resistance of 10 Ω is assumed.

The collector voltage is set to 5 volts, with no source series resistance. The base voltage is set to sweep from 0.7 volts to 1.4 volts in seven 1-second time steps, again with no source series resistance. Fig. 3 in Section 3.15 illustrates the connection and polarity of these two external sources.

Esc back to the Program Control menu if necessary, move the cursor to the top of the menu, and press Enter to proceed with the solution. The solution will proceed to solve for equilibrium and the initial steady-state condition. Convergence to the steady-state condition is slow for this problem, with many iterations required. The base voltage will then begin ramping from 0.7 to 1.4 volts in 0.1-volt increments. Don't hesitate to press Enter during the course of the solution if you want to examine the results for a particular bias point. At the completion of the solution, press Enter to get the Graph Menu. Select the User-Defined option and set up the graph by choosing Base Voltage for x, Base Current for y1, and Collector Current for y2. Change the y scaling to logarithmic. When you have made these entries, press Esc to exit the setup menu and graph the results. It should appear similar to Fig. 16. Because both I_b and I_c are on a logarithmic scale, the spacing between the two curves is a direct indication of the transistor gain, β . The collector current climbs more quickly than the base current because the collector current has an ideality

factor close to 1, while the base current is due mostly to recombination in the base-emitter-junction space-charge layer and therefore has an ideality factor close to 2. The curves bend to the right at high current levels because of the voltage drop across the internal base series resistance.

Fig. 16 Gummel plot displaying transistor gain.

The current gain is low when the base drive is small, and increases to a maximum of about 200. Examine the other plots available, remembering that only the last four operating points are available for detailed views of the spatially-dependent values. When finished, Esc back to the Program Control menu.

To generate a common-emitter characteristic curve for this device, both the base and collector voltage sources must be modified. First move the cursor to the Circuit line, press Enter, then select the Base Contact. Keep the base connected, but enter a source resistor of 10 k Ω (1E4 ohms). Esc back to the Program Control menu. Move the cursor to the Base Voltage line, press Enter, and select a Fixed voltage of 11 V. This source voltage, minus the ≈ 1 -volt base-emitter voltage, will force 1 mA to flow through the 10 k Ω source resistor. Move the cursor to the Collector Voltage line, press Enter, and select a Sweep. Set the initial voltage to 0.1 volt and the final voltage to 0.8 volt. Esc back to the Program Control menu and Proceed with the solution.

This problem is also a challenge for convergence, so it will take some time. When the seven steps are completed, press Enter to get the Graph Menu and select the User-Defined graph option. Set up the graph so that x is the Collector Voltage, y1 is the Collector Current, and both axes are scaled linearly. Fig. 17 illustrates the resulting graph.

Fig. 17 Common-emitter characteristic curve.

If you have worked through these example problems, it should be apparent that the possibilities for variation are limitless. How does the substrate doping level affect the onset of an inversion layer? What is the effect of different junction depths on diode behavior? How fast does a solar cell's short-circuit current decay after the light shuts off? What happens when you shine light on a transistor? What about double heterostructures? What about the effect of temperature? Hopefully, PC-1D will raise as many questions for you as it answers.

Provided on the DRIVERS distribution diskette is a program called PC1DPLOT.EXE. This program can plot on the screen and on a compatible plotter data that has been stored in a properly-formatted ASCII file. This includes all plot data saved from within PC-1D (having a filename of the form *.PDF). You can edit the plot-data file prior to plotting it, allowing you to remove data points, change the labels, or adjust the axis scaling.

If you plan to use PC1DPLOT to generate a hard-copy plot, then before invoking PC1DPLOT, make certain that the plotter driver (VDIPLOT.RSYS) has been loaded by your CONFIG.SYS, that the correct COM port has been selected (if other than COM1), and that the COM port has been configured for the correct baud rate.

PC1DPLOT can be invoked simply by typing PC1DPLOT in response to the DOS prompt. You will be asked to provide the name of the data file. If you don't give a file extension, .PDF is assumed. If you have DOS version 3.0 or later, you can also give the filename directly on the command line:

```
C> PC1DPLOT myfile1
```

In either case, if the data file specified can't be found, or if it contains a format error, an error message will be displayed and the program terminated. Otherwise, you will soon see a plot of the data on the screen, which you may choose to send to the pen plotter. Note that, depending on which display you have (CGA, EGA, or VGA), the plot generated by the pen plotter may be slightly different than what you observe on the screen. PC1DPLOT will take advantage of the resolution of the plotter to generate more precise axis scaling than may have been possible on the display. The data itself will be the same.

The format of the plot-data file is illustrated in Table 5. The comments in brackets are explanatory only and should not appear in the file. Although PC-1D generates a maximum of four data sets for a single plot, files you create yourself can have up to seven data sets, each displayed using a different linetype. Each data set is represented by a series of x-y points. Each x-y pair appears together on the same line in the file, separated by one or more spaces. These points will be connected by straight line segments in the order they appear in the file. Otherwise, there are no restrictions on the ordering of the points; you can even backtrack to trace out a circle if you want. You are limited to maximum of 8000 x-y points for all data sets combined. You need not have the same number of points in each data set. The only place that blank lines are allowed in the file are the first six lines. In particular, make sure that a blank line is not accidentally added to the end of the file.

Table 5
PC1DPLOT Data File Format (*.PDF)

Doping Concentrations	{Line 1: title of graph}
Distance from Front (#m)	{Line 2: x-axis label}
cm-3	{Line 3: y-axis label}
PC-1D	{Line 4: first bottom comment}
22 JAN 89 12:40p	{Line 5: second bottom comment}
	{Line 6: third bottom comment}
0.00000E+00 1.00000E-05 0	{Line 7: x_{min} , x_{max} , x-scaling}
1.00000E+00 1.00000E+20 1	{Line 8: y_{min} , y_{max} , y-scaling}
6	{Line 9: number of points in first data set}
0.00000E+00 1.00000E+20	{ x-y pairs for first data set}
1.65907E-07 6.43778E+19	
6.93163E-07 4.58479E+16	
1.22042E-06 4.47132E+09	
3.34036E-06 1.00000E+00	
1.00000E-05 1.00000E+00	
7	{ number of points in second data set}
0.00000E+00 1.00000E+16	{ x-y pairs for second data set}
1.65907E-07 1.00000E+16	
6.93163E-07 1.00000E+16	
1.22042E-06 1.00000E+16	
2.93532E-06 1.00000E+16	
6.17568E-06 1.00000E+16	
1.00000E-05 1.00000E+16	

Several items require further explanation. You can put the # character in either or both of the x- and y-axis labels. If you do so, then the # sign will be replaced with an appropriate metric SI prefix. This assumes that the corresponding data has been scaled to the base SI unit (i.e. meters, not centimeters). All titles are limited to 31 columns or less, comments are limited to 63. Lines 7 and 8 control the axis scaling. The min and max values define the range that the axis must cover. These are normally the min and max values of data, but this need not be the case. For example, you may want to consistently plot quantum efficiency on a scale of 0 to 100%, even if the data falls in a narrower range. You can set the min and max values both to zero, in which case PC1DPLOT will automatically scale the graph to fit the data. The x- and y-scaling values are either 0 for linear scaling or 1 for logarithmic scaling.

Finally, if you place a second number on the line following the number indicating the number of items in a data set, then the data set which follows will be plotted as a series of individual markers not connected by lines. The second number is interpreted as the size to make the markers (in percent of full screen). By repeating the same data set twice, once with the marker size value and again without, you can generate a plot that has both markers and connecting lines.

PC-1D Program Development:

Basore, P.A., "PC-1D Version 3: Improved Speed and Convergence," 22nd IEEE Photovoltaic Specialists Conf., (October 1991).

Basore, Rover, and Smith, "PC-1D Version 2: Enhanced Numerical Solar Cell Modeling," 20th IEEE Photovoltaic Specialists Conf., (September 1988).

Basore, P.A., "Numerical Analysis of Semiconductor Devices on Personal Computers," Proceedings of the Sixth Biennial University/ Government/ Industry Microelectronics Symposium, pp. 73-76 (1985).

International Business Machines Corporation, IBM Personal Computer Graphics Development Toolkit, IBM, Boca Raton (April 1987).

Microsoft Corporation, IBM Personal Computer Macro Assembler Version 2.00, International Business Machines Corporation, Boca Raton (August 1984).

Microsoft Corporation, IBM Personal Computer Pascal Compiler Version 2.00, International Business Machines Corporation, Boca Raton (February 1984).

Rover, Basore, and Thorson, "Solar Cell Modeling on Personal Computers," 18th IEEE Photovoltaic Specialists Conf., Las Vegas (October 1985).

Physical Models:

Adachi, J. Appl. Phys. 58, p. 1 (1985). (GaAs, AlGaAs electronic properties)

Arora, Hauser, and Roulston, IEEE Trans. on Electron Devices ED-29(2), p. 292 (February 1982). (Silicon transport properties)

Aspnes, Kelso, Logan, and Bhat, J. Appl. Phys. 60, p. 754 (1986). (GaAs, AlGaAs optical properties)

Aspnes and Studna, Physical Review B 27, p. 985 (1983). (InP, Ge optical properties)

ASTM Standard E891-R4, Annual Book of ASTM Standards (1985) (AM1.5 Direct Spectrum)

ASTM Standard E892-R4, Annual Book of ASTM Standards (1985) (AM1.5 Global Spectrum)

Camassel and Auvergne, Phys. Rev. B 12, p. 3258 (1975). (Ge electronic properties)

Camassel, Auvergne, and Mathieu, J. Appl. Phys. 46, p. 2683 (1975). (GaAs electronic properties)

del Alamo, Swirhun, and Swanson, Solid-State Electronics 28 (1-2), p. 47 (1985). (Silicon transport and recombination properties)

Fonash, J., Solar Cell Device Physics, Academic Press, N.Y. (1981). (Ge electronic properties)

Haas, J. Phys. Chem. Solids 23, p. 821 (1962). (Ge optical properties)

Haug, J. Phys. C: Solid State Phys. 16, p. 4159 (1983). (GaAs, AlGaAs recombination coefficients)

Karpova and Kalashnikov, Proceedings of the International Conference on the Physics of Semiconductors, p. 880 (1962) (Ge recombination coefficients)

Klausmeier-Brown, Lundstrom, and Melloch, Appl Phys. Lett. 52 (26), p. 2255 (June 1988). (GaAs electronic properties)

McLean, T.P., Progress in Semiconductors, Vol. 5, Heywood, London (1960). (Ge electronic properties)

Morin, Phys. Rev. 94, p. 1525 (1954). (Ge transport properties)

Prince, Phys. Rev. 92, p. 681 (1953). (Ge transport properties)

Rajkanan, Singh, and Shewchun, Solid-State Electronics 22, p. 793 (1979). (Silicon optical properties)

Sigai, Abrahams, and Blanc, J. Electrochem. Soc 119, p. 952 (1972). (AlGaAs transport properties)

Sproul, Green, and Zhao, Appl. Phys. Lett. 57 (3), p. 255 (July 1990). (Intrinsic concentration in silicon).

Swanson and Swirhun, "Characterization of Majority and Minority Carrier Transport in Heavily Doped Silicon," Sandia National Laboratories contractor report, SAND87-7019 (November 1987). (Silicon transport and recombination properties).

Sze, S., Physics of Semiconductor Devices, 2nd. Ed., J. Wiley and Sons, New York, 1981, p. 850. (Si electronic properties)

Takeshima, J. Appl. Phys. 58, p. 3846 (1985). (AlGaAs recombination coefficients)

Takeshima, Jap. J. Appl. Phys. 22, p. 491 (1983) (InP recombination coefficients)

Thekaekara, Solar Energy Engineering, Ed. by A.A.M. Sayigh, Academic Press, New York, 1977, p. 40. (AMO Spectrum)

Thurmond, J. Electrochemical Soc. 122, p. 1133 (1975). (Si, GaAs, Ge optical properties)

Turner, Reese, and Pettit, Phys. Rev. A 136, p. 1467 (1964) (InP electronic properties)

Tyler and Woodbury, Phys. Rev. 102, p. 647 (1956). (Ge transport properties)

Vilms and Garrett, Solid State Electron. 15, p. 443 (1972). (GaAs transport properties)

Walukiewicz, et. al., J. Appl. Phys. 51, p. 2659 (1980). (InP transport properties)

Wang and Neugroschel, IEEE Electron Device Lett. 11 (12) (December 1990). (Silicon minority hole mobility).

Wiley, J., Semiconductors and Semimetals, Vol. 10, Academic Press, N.Y. (1975). (InP, GaAs, AlGaAs transport properties)

Yahia and Coutts, "Uncertainties in the physical parameters used in modeling InP solar cells", to be published in Solar Cells (InP electronic properties)

Yahia, Wanlass, and Coutts, 20th IEEE Photovoltaic Specialists Conf. (September 1988) (InP recombination coefficients)

Numerical Method:

Basore, P.A. "Numerical Modeling of Textured Silicon Solar Cells Using PC-1D," IEEE Trans. on Electron Devices, ED-37 (2), (February 1990).

Basore, P.A., "Essential Elements in the Numerical Modeling of Solar Cells," 4th Photovoltaic Science and Engineering Conf., (February 1989).

Grossman and Hargrove, "Numerical Solution of the Semiconductor Transport Equations with Current Boundary Conditions," IEEE Trans. on Electron Devices ED-30 (9), pp. 1092-1096 (September 1983).

Dongarra, Moler, Bunch, and Stewart, LINPACK User's Guide, SIAM, Philadelphia (1979).

Lundstrom, Schwartz, and Gray, "Transport Equations for the Analysis of Heavily Doped Semiconductor Devices," Solid-State Electronics 24, pp. 195-202.

Lundstrom and Schuelke, "Numerical Analysis of Heterostructure Semiconductor Devices," IEEE Trans. on Electron Devices ED-30 (9), pp. 1151-1158 (September 1983).

Mock, Analysis of Mathematical Models of Semiconductor Devices, Boole Press, Dublin.

Selberherr, Analysis and Simulation of Semiconductor Devices, Springer-Verlag, New York.

Sutherland and Hauser, "A Computer Analysis of Heterojunction and Graded Composition Solar Cells," IEEE Trans. on Electron Devices ED-24, pp. 363-372 (1972).

Thorson, G.M., "A Solar Cell Simulator for Personal Computers," M.S. Thesis, Iowa State University, Ames, IA, 1985.

DISTRIBUTION:

1	6200	B. W. Marshall
1	6220	D. G. Schueler
1	6220	A. Van Arsdall
20	6220	J. Pineau, Solar Library
1	6224	D. E. Hasti
10	6224	P. A. Basore
1	6224	J. M. Gee
1	6224	W. K. Schubert
1	6224	D. S. Ruby
5	3141	S. A. Landenberger
1	3145	Document Processing For DOE/OSTI
3	3151	G. L. Esch, Acting
1	8523	R. C. Christman

Sandia National Laboratories