

PC1D Help Index

What's new in Version 5.9?

Getting Started

Using the Device Schematic

Numerical Method

Convergence and convergence failure

Physical Constants

Bibliography

Commands

File menu

Device menu

Excitation menu

Compute menu

Graph menu

View menu

Options menu

Help menu

Program Description

PC1D is a computer program written for IBM-compatible personal computers which solves the fully coupled nonlinear equations for the quasi-one-dimensional transport of electrons and holes in crystalline semiconductor devices, with emphasis on photovoltaic devices. This version of the program is supported and distributed by the Photovoltaics Special Research Centre at the University of New South Wales in Sydney, Australia 2052. Only licensed copies of the program are authorized for use. A licensed copy may be loaded and used on multiple computers or on a network provided the licensee maintains records of the number and location of these authorized copies and can ensure that all notifications and updates are distributed to everyone using these copies. Licensed copies are available from the PV Centre for a fee of AU\$150, which may be paid by credit card or by cheque. Contact the PV Centre via e-mail at pc1d@unsw.edu.au to receive an order form via fax.

PC1D runs under Windows 95/98/ME/XP/NT, and requires at least an 80386 CPU and an 80387 math coprocessor (note that most 80486 and Pentium processors have the math coprocessor built-in). It will also run on Windows 3.1 if Win32s is installed (Win32s 1.71 or later is required).

Only one file is necessary to run the program, PC1D.EXE. The additional file PC1D.HLP provides on-screen help, and several additional files are provided which contain material parameters for selected semiconductors, standardized solar spectra, and example problems. All of the files can be simply copied into the directory of choice; no setup program is required. To store files of different type in different directories, see the instructions for the Options menu.

PC1D will continue to be improved and your suggestions are appreciated. Submit them via e-mail at the address shown above. Of particular interest are any computation errors that may arise, and improved values for material parameters as they become available. Licensed users will receive update notices. Those who register an e-mail contact address will be provided with maintenance updates of the program and its associated files via e-mail at no cost.

Getting Started

Using PC1D is a three step process:

1. Set up the simulation parameters. This includes the device and material parameters, and the excitation to be applied to the device.
2. Run the simulation.
3. Examine the results.

PC1D has three different types of displays (called *views*), each useful for a different step of this process.

When setting up the simulation parameters, you will want to use the *Parameter View*. This gives a list of simulation parameters, and a schematic diagram of the device being simulated, giving you visual feedback when parameters are changed. This view also gives a shortcut: you can double-click on any line to bring up a dialog box allowing you change that parameter. You can also double-click on elements of the diagram to change them.

While you run a simulation, you can switch to the *Four-Graph View*. This fills the screen with four graphs of quantities that are commonly of interest (for example, the carrier velocities; generation and recombination; etc). You can watch these quantities change as the program advances in the solution of the problem.

The *Interactive Graph View* is designed for intensive study of a particular graph. It allows you to zoom into regions of the graph which interest you; examine the values of individual points; and copy the graph values to another Windows program (e.g. a spreadsheet) for further analysis.

But, PC1D is designed to be flexible! You can switch between views at any time (even while a simulation is running).

If you are just using PC1D for Windows for the first time, you should take a look at the example parameter files supplied with the program. Spend a little time getting used to the various views, and the methods of changing parameters. A few common questions are answered below.

How do I change the parameters for a simulation?

Either:

- Use the Display menu and Excitation menu; or
- Double-click on the parameters name in the Parameter view (This is the textual list of parameters which you see when you first run PC1D); or
- Double-click on the device schematic to change device parameters. A dialog box will appear, allowing you to specify values for the parameters which you are interested in.

How do I get PC1D to calculate the results? (How do I run a simulation?)

Either:

- Use the Compute menu and select Run; or
- Press the Run button in the toolbar.

In the old DOS versions of PC1D, the program displayed four graphs while it was calculating the results. How can I make this happen in the Windows version?

Switch to Four Graphs view. You can do this from the View menu, or by pressing the FourGraphs button in the toolbar.

0 How can I examine the results of a simulation?

- 1 Switch to interactive graph view. You can do this using the Graph menu.
- 2 You can zoom into any region of the graph using a mouse or keyboard. Values from the graph can be copied to the clipboard and imported into another program, e.g. a spreadsheet.
- 3 Or, if you are only interested in the values of Voc, Isc and Pmax, they are displayed in the parameter view under the heading ***** RESULTS *****

Convergence and 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 non-equilibrium solution difficult. Several measures have been taken within PC1D to assist convergence, but despite this you will eventually generate problems that will not converge. This section offers some advice on how to avoid convergence failure, and how to deal with it when it does occur.

Dynamic renoding

This feature was introduced in PC1D 4.2, and reduces the likelihood of non-convergence. The basic idea is that when PC1D detects that part of the device is not converging well, it increases the number of finite elements in the difficult area. This helps to ensure that the assumptions which PC1D makes about the behaviour of solution variables remain valid.

Dynamic renoding is particularly significant for problems involving reverse biased current sources.

In general, convergence failure occurs either because:

- (a) the dynamic renoder ran out of nodes (currently, there is a limit of 500);
- or (b) the solution is trying to reach a final state that is too far removed from the initial state;
- or (c) the situation is unphysical and has no solution.

Situations which are known to interfere with convergence

The following situations should be avoided as they make convergence difficult:

1. Locating an electrical contact in a region that, in equilibrium, is either very lightly doped or depleted. If the region in equilibrium is clearly of one type or the other, that type will be assumed to be the polarity of the contact, even if subsequent excitation causes the carrier concentrations at that point to become inverted.
2. Connecting a shunt element between two dopant regions that are both isolated by a junction from an electrical contact. This situation occurs in modeling series-connected multijunction devices. It is generally best to model these devices instead as three-terminal devices, then infer the two-terminal behavior from the three-terminal results.
3. Applying a large forward voltage with no current-limiting resistor. The currents in the device can become huge in this case, and numerical overflow can occur.
4. Low velocity saturation. This is a problem when carriers are trying to go faster than the velocity limit. This problem can be avoided by setting the limit to zero (which disables velocity saturation), or by choosing fixed mobility rather than variable. The problem is even more severe if total velocity saturation is specified for the numerical method.

Improving convergence by controlling the numerical method

The following actions can be taken to try to get the problem to converge without changing the definition of the problem.

1. Change the element size factor in the Compute:Numerical dialog box.

By decreasing the element size factor, you increase the number of elements in the problem. Decreasing this value will tend to improve convergence, until the maximum of 500 elements is reached. However, the problem will take longer to solve. Sometimes, you will encounter situations where convergence can be improved by *increasing* the element size factor.

2. Adjust the normalized potential clamp (also in the Compute:Numerical dialog box).

A smaller value (between 0.1 and 1) will sometimes improve convergence, although some problems benefit from larger clamp values (5-10). It is particularly helpful to reduce this value when voltage, current or light are applied abruptly. Small clamp values increase the time required to solve the problem, especially problems where large reverse-bias voltages are applied.

3. Ensure that *Psi* and *Phi clamping* (in the Compute:Numerical dialog box) are not both disabled.

4. Turn off *Total velocity saturation* (in the Compute:Numerical dialog box) unless you need to include this effect (as you might for some heterostructures where current is limited by velocity saturation in areas of sharp carrier-concentration gradients).

5. On some occasions the difficulty simply may be that the problem converges very slowly. In this case, you should increase the time limit in the Compute:Numerical dialog box.

Improving convergence by imposing excitation gradually

This is the most effective way of improving convergence. For example, to solve a silicon junction forward-biased to 0.8 volts, you may need to perform an interim solution at 0.6 volts first. You should solve the problem for steady-state at 0.6 volts, then change the bias to 0.8 volts and use Compute:Continue. This will solve for steady-state, using the

interim solution as the starting point for the next solution.

Similarly, it may be necessary to increase the light intensity in steps, say 10 mW/cm², 100 mW/cm², then 1 W/cm².

If you have a particular interest in understanding exactly why a particular problem did not converge, you can enable *Graphs after every iteration* in the Compute:Numerical dialog, and create a user-defined graph of *Convergence Error* (plotted on a log scale) versus *Distance from front*. This will show how far each element of the device is from convergence.

What's new in Version 5.9?

(Release date: June 03).

Graphs of minority carrier lifetime and diffusion length should now work properly in equilibrium. In the past, these graphs suffered from a loss of precision, making them jagged.

The short description of some batch parameters (related to doping) were changed in the QuickBatch dialog to improve clarity.

What's new in Version 5.8?

(Release date: Dec 02).

A bug in the velocity saturation code was fixed. This meant that the mobility reduction at high fields was being overestimated by a factor of about two. This bug did not apply when 'total velocity saturation' was selected.

What's new in Version 5.7?

(Release date: July 02).

One minor new feature has been added, which will only be of interest to programmers.

There is a new command-line option /g which allows PC1D to be run from an external program.

PC1D /g filename.prm opens the parameter file 'filename.prm' silently, runs the simulation, copies the contents of the interactive graph to the clipboard, and then exits. Note that if the PRM file had a one-line batch file associated with it, you could modify the batch file before running PC1D in order to change a model parameter.

Changes introduced in Version 5.6

(Release date: Sept 01).

Rear surface texturing now works. A bug in all previous versions was causing rear surface texturing to be applied at the front instead of the rear.

A new graphable function, 'Pri-Surface Total reflectance' has been added, to make it easier to match experimental reflectance data. The Quantum Efficiency graph now displays this total reflectance, instead of separate curves for front reflectance and escape.

BUGFIX: The batch commands for bandgaps were muddled up (Bandgap, AbsEd1, AbsEd2, AbsEi1, AbsEi2).

Changes introduced in Version 5.5

(Release date: Aug 00).

An error in the diffusion length graph has been fixed, and the limitations on mobility have been relaxed, allowing simulations of very low mobility cells. Some minor typographical errors in the numerical method section of the help file have been corrected.

In rare circumstances, previous versions of PC1D 5 would crash when exiting. This problem has been fixed.

Changes introduced in Version 5.4

(Release date: June 00).

- BUGFIX: Under certain circumstances, quantum efficiency calculations would not converge at long wavelengths. (This was due to a compiler bug, and only occurred in PC1D 5.3). This problem has been fixed.
- Graphs now look better when the PC1D window is only occupying part of your screen. The graphs remain readable down to small window sizes.
- A 'reflectance' shortcut button has been added to the toolbar, as a convenient way of adjusting the device reflectance.

No simulation results are affected by any of these changes.

New features introduced in Version 5.3 - Major batch enhancements

Internal Batch mode ('QuickBatch' button on toolbar)

A batch run can now be generated from inside PC1D, without having to use an external spreadsheet program. Just specify which parameters you want to vary, and over what range, and PC1D will do the rest.

Graphs saved between simulations

PC1D now saves graphs from the last 100 simulations. You can access them by pressing PageUp and PageDown in the Interactive Graph View. This is especially useful for batch runs: if you run a batch where only one parameter is varying, by pressing PgUp and PgDn, you can graphically see the effect of varying that parameter.

Increased speed (again!) – Now twice as fast as version 5.0 !

With the new models turned off, PC1D is now five times as fast as version 4.6, or seventeen times as fast as the old DOS version (PC-1D 3.3). With the new models turned on, it is fourteen times faster than the DOS version.

Parameter view displays more information

Values for recombination, applied circuit, light intensity, series resistance, etc are now displayed in the parameter screen. This will help ensure that you are performing the simulation that you intended. It also uses subscripts and superscripts to aid readability.

Slight convergence improvement

Some poorly-behaved problems will now converge, due to improved numerical precision in some of PC1D's internal functions.

Minor features

- Light intensity limit increased: Some characterisation techniques using laser pulses result in extremely high light intensity. These situations can now be simulated.
- If you stop a simulation, change the light sources, and continue the simulation, PC1D will now recalculate the photogeneration before continuing.
- New graphable function, 'IQE adjusted for light bias', to easily obtain the quantum efficiency of a device which is being illuminated by a secondary (constant) light source.
- New shortcut batch parameters (BulkTau, FrS, RrS, FrIntRefl, RrIntRefl).

Changes introduced in Version 5.2

(Release date: Sept 98).

Bugfix - deleting batch files could cause PC1D to crash

Previous versions of PC1D didn't check whether batch files had been deleted. PC1D would crash at the end of a simulation if the batch file had been deleted. This problem has now been fixed.

Bugfix - batch files with many parameters didn't always work

Version 5.1 didn't always work properly for batch files with more than about 18 parameters. It will now work properly with up to 30 parameters.

Better use of exponential notation

- Dialog boxes now use exponential notation for large values, instead of (cumbersome) fixed-point notation. For example, 43000000000 is now displayed as 4.3e10. This is particularly helpful for parameters such as intrinsic carrier concentration.
- The range for which scientific notation (rather than fixed point) is used can be changed from a new Options dialog box in the Options menu. It will affect dialog boxes, as well as the parameter view.

Numerical method described in help files

The numerical method used by PC1D is now described in detail in the help files.

Changes introduced in Version 5.1

(Release date: Mar 98).

Parameter View font selection

- The font to be used for the parameter view can now be selected (select *Font* in the *Options* menu). If you have a large screen, you may want to select a smaller font.

Better batch files

- There is now no intrinsic limit to the length of a batch file. Previously, batch files were limited to 100 lines. The maximum number of lines now depends on the operating system you are using. In Windows 95, batch files are limited to a few hundred lines (more if smaller font sizes are used in Parameter View). This limitation doesn't apply to Windows NT, and Microsoft will probably fix the problem for Windows 98.
- Batch files can now have 30 fields across (although only the first 12 will print on A4 paper, unless you select a very small font size or copy the results into a spreadsheet program).

Increased speed (again!)

- Simulation speed is now three times as fast as PC1D 4.6, or eight times as fast as PC1D 3.6. On a 200Mhz Pentium, an IV curve for a simple cell now takes slightly less than 1 second.

Non-convergent problems

- PC1D is now much better at detecting non-convergent problems. If it detects a problem it can't solve, it will stop immediately. It won't display bad numerics in graphs.
- The annoying 1 second delay when trying to stop a non-convergent problem has been dramatically reduced.

Minor bug fixes

- Copying batch file data to clipboard would occasionally add a line of garbage at the end. This has been fixed.
- The program was always disabling rear external photogeneration files when reloading. It now saves and loads correctly.
- The mouse didn't work well with non-simple curves in interactive graphs. For example, if you were sweeping voltage from +ve to -ve, it was not possible to select points with the mouse. Also, Voc, Isc and Pmax weren't always displayed. These problems have now been fixed.

Changes introduced in Version 5.0

Version 5 is the first 32-bit (Windows 95 / Windows NT) version of PC1D. It will also run on Windows 3.1 if Win32s is installed (Win32s 1.71 or later is required). It was released on 8 Sept 97. The new features are:

Increased speed

- Simulation speed is now twice as fast!

Trap-assisted tunnelling

- Trap-assisted tunnelling can now be modelled using the Hurkx model for field-enhanced recombination. Access it through the Device:Material:Recombination dialog box.

Experimental data graphs

- External files of experimental data can be displayed simultaneously with simulation results. This makes it much simpler to fit simulation parameters to experimental results.

Increased simulation domain

- Extra batch parameters for shunt elements, contact positions, bandgap, and intrinsic concentration.
- Internal reflection can now be set to 100% (it was previously limited to 99%).

Better graphs

- All graphs can now be user-defined. You can redefine any of the existing graphs by (a) using the Define command in the Graph menu; or (b) double-clicking in the border area of any interactive graph.
- The graphs you select in the Four-graphs view and Interactive graph view are now saved with the excitation and parameter files. This is particularly helpful for IQE scans, etc, because it means you don't have to redefine your graphs every time you restart PC1D. It also saves your auxiliary and experimental data graphs.
- In interactive graph view, press CTRL+arrow keys to move the caret rapidly.

User interface enhancements

- Support for long filenames.
- Minor features such as ToolTips for toolbar buttons.
- Tabbed dialog boxes! Many of the old dialog boxes were combined into tabbed dialogs, so the menu structure is now much simpler and less daunting. This is a particularly big improvement for reflectance, and for light sources.
- The parameter view no longer flickers while running a simulation. (This is one contributor to the increased speed).

Backwards compatibility

- Loads files created by any previous version of PC1D
- Option to save material, device, excitation and parameter files in a form that allows them to be read by PC1D 4.5 (Of course, any new features will be lost when saved in the old version). This option appears in all of the "Save As" dialog boxes.

Minor changes and bug fixes

- The program now gives correct results for rear illumination of a device made from more than one material.
- External files of absorption data are now interpolated logarithmically rather than linearly.

Changes introduced in Version 4.6

Version 4.6 was a very minor maintenance release, fixing an error which occurred when n_i was small, such as with large bandgap materials at low temperature. In such cases, the program erroneously introduced a large resistance at the contacts.

Changes introduced in Version 4.5

Version 4.5 was a minor maintenance release, fixing a few problems which were discovered after the release of 4.4:

- The list of recently used files under the File menu now works properly.

- The program now gives correct results when monochromatic light is used with zero intensity (such as happens when simulating the transient response to a laser pulse).
- The default values for free-carrier absorption in silicon now match the values recommended in the help file.
- A few more parameters can now be used in batch files: shunt elements, bandgap and intrinsic concentration.
- The file locations “Problem Parameters” external file path was removed, as it was not used by the program. The initial location for .PRM files should be set using the standard Windows “Properties” function (in Win 3.1, select the icon and press SHIFT+Enter).

Changes introduced in Version 4.4

4.4 was a minor maintenance release (Release date: Dec 96).

The silicon material parameters were updated to be consistent.

Also, a few minor bugs were fixed. E.g. in the device schematic, diodes are now displayed the correct way around.

New features introduced in Version 4.3 - Device diagram, free-carrier absorption, and total velocity saturation.

Version 4.3 was a major release (Release date: Oct 96).

In addition to minor bug fixes and improvements to the online help, the following new features were incorporated into PC1D Version 4.3:

User interface enhancements

- A diagram of the device is now displayed in the Parameter View, providing visual feedback of doping, texturing, and internal shunt elements. This makes it much easier to recognize mistakes when designing complex devices. You can also double-click on parts of the diagram to change them, which gives another method for setting up parameters.
- The behaviour of the scroll bars has been improved.

Increased capacity

- The maximum number of timesteps has been increased to 200.

Physics

- The majority-carrier mobility model for silicon has been revised to be more accurate in the vicinity of room temperature. The model now agrees better with the 1981 data of Thurber et al.
- The band-to-band recombination coefficient for silicon was corrected from $9.5E-14$ to $9.5E-15$ cm³/s.
- The ratio of N_c/N_v for silicon was changed from 2.8 to 1.06 to be consistent with the “new” lower value of intrinsic carrier concentration of $1E10$ cm⁻³ at 300K which is now in common use.
- Free-carrier absorption can now be modelled. This improves the accuracy of simulations of heavily-doped devices.
- Velocity saturation can now be modelled accurately. Previous versions only limited carrier velocity due to a high electric field. You can now limit the velocity due to both drift and diffusion. To do so, turn on *Total velocity saturation* in the Compute:Numerical dialog box. Note, however, that invoking this feature will significantly slow down your solution and is only recommended when this effect is important to your device. Normally, it will only affect heterostructures where the current is limited by thermionic emission over an energy barrier.

New features introduced in Version 4.2

(Release date: Aug 96).

Improved convergence

- Dynamic renoding was introduced into the solution code. This improves convergence of many problems, especially reverse bias and floating junctions.
- The internal equations were changed back to the old ones used in Version 3. It turns out that the old equations have superior convergence properties.
- 8 extra plot functions were added.

Increased speed

- Photogeneration now 4 times faster. (Makes steady-state problems 25% faster overall).

New features introduced in Version 4.1

4.1 was a minor maintenance release (Release date: July 96).

Several bugs were fixed. In addition:

- You can now use a previously solved solution as a starting point for a new one, using the ‘Continue’ feature in the Compute menu.
- Width of batch files was increased to 12 parameters.
- A new plot function (Convergence Error) was added.

Version 4.0 – Initial Windows release

Version 4.0 was the first version of PC1D for Windows. It was showcased at the 25th IEEE Photovoltaics Specialist Conference in Washington DC, and publicly released in June 96. It replaced the DOS version, PC-1D version 3.3.

File menu commands

The File menu offers the following commands:

Description	Enter or modify a text description of the contents of this parameter file.
New	Creates a new parameter file.
Open	Opens an existing parameter file.
Close	Closes an opened parameter file.
Save	Saves an opened parameter file using the same file name.
Save As	Saves an opened parameter file to a specified file name.
Print	Prints the current window.
Print Preview	Displays the current window on the screen as it would appear printed.
Print Setup	Selects a printer and printer connection.
Exit	Exits PC1D.

View menu commands

The View menu offers the following commands:

Toolbar	Shows or hides the toolbar.
Status Bar	Shows or hides the status bar.
Parameter View	The Parameter view is the start-up default. This text-type screen lists all of the parameters that define your problem, and at the end there is a section for reporting the results of calculations. You can double-click the mouse on most lines in the Parameter view to open a dialog box to modify that parameter. A schematic diagram of the device is also displayed, providing visual feedback of doping, texturing, and internal shunt elements. This makes it easier to recognize mistakes when designing complex devices. You can also double-click on parts of the diagram to change various parameters.
Four-Graph View	The Four-Graph view provides a quick overview of what's happening in your device. Double click on any of the four graphs to select it for detailed examination using the Interactive-Graph view.
Interactive-Graph View	The Interactive Graph view allows you to examine specific data values on a graph. If the data you need is not in one of the predefined graphs, you can create a user-defined graph using the Graph menu to access any of 55 different functions.

Interactive Graph View

This is the view you will want to use most often when inspecting solution results. It displays a single graph, and allows you to zoom into regions of the graph, and extract the values of individual points.

To zoom into part of the graph:

With the mouse: Press the left mouse button. Drag over the area you want to zoom into. Release the left mouse button. (To cancel the zoom, press the right mouse button or the ESC key).

With the keyboard: Use the arrow keys to move the caret (the blinking line) to one end of the area you are interested in. Press ENTER. Move the caret to the other end. Press ENTER.

To zoom out:

With the mouse: Press the right mouse button.

With the keyboard: Press the ESC key.

Zooming out when the graph is already fully 'zoomed out' will return you to the four-graphs view.

To examine graphs from previous simulations:

Press PAGEUP or PAGEDOWN to view the graphs from the last 16 simulations.

Help menu commands

The Help menu offers the following commands, which provide you assistance with this application:

- Index Offers you an index to topics on which you can get help.
- Using Help Provides general instructions on using help.
- About Displays the version number of this application.

New command (File menu)

Use this command to create a new parameter file in PC1D.

You can open an existing parameter file with the Open command.

Shortcuts

Toolbar:



Keys:

CTRL+N

Open command (File menu)

Use this command to open an existing parameter file. The parameters for a problem, including a description of both the device and excitation, are stored in a binary-encoded parameter file with a PRM suffix. The parameter file also specifies external data files that may be necessary to describe parameters that are a function of position, wavelength, or time.

You can create new parameter files with the New command.

Shortcuts

Toolbar:	
Keys:	CTRL+O

File Open dialog box

The following options allow you to specify which file to open:

File Name

Type or select the filename you want to open. This box lists files with the extension you select in the List Files of Type box.

List Files of Type

Select the type of file you want to open:

Drives

Select the drive in which PC1D stores the file that you want to open.

Directories

Select the directory in which PC1D stores the file that you want to open.

Close command (File menu)

Use this command to close all windows containing the active parameter file. PC1D suggests that you save changes to your parameter file before you close it. If you close a parameter file without saving, you lose all changes made since the last time you saved it. Before closing an untitled parameter file, PC1D displays the Save As dialog box and suggests that you name and save the parameter file.

Save command (File menu)

Use this command to save the active parameter file to its current name and directory. When you save a parameter file for the first time, PC1D displays the Save As dialog box so you can name your parameter file. If you want to change the name and directory of an existing parameter file before you save it, choose the Save As command. The parameters for a problem, including a description of both the device and excitation, are stored in a binary-encoded file with a PRM suffix. The parameter file also specifies external data files that may be necessary to describe parameters that are a function of position, wavelength, or time.

Shortcuts

Toolbar:



Keys:

CTRL+S

Save As command (File menu)

Use this command to save and name the active parameter file. PC1D displays the Save As dialog box so you can name your parameter file.

To save a parameter file with its existing name and directory, use the Save command.

File Save As dialog box

The following options allow you to specify the name and location of the file you're about to save:

File Name

Type a new filename to save a parameter file with a different name. PC1D automatically adds the extension you specify in the Save File As Type box, if you don't supply a different one.

Save File As Type

Choose between the latest version of PC1D, or an old file format which can be opened by older releases of PC1D. Of course, new features can't be saved in the old format.

Drives

Select the drive in which you want to store the parameter file.

Directories

Select the directory in which you want to store the parameter file.

1, 2, 3, 4 command (File menu)

Use the numbers and filenames listed at the bottom of the File menu to open the last four parameter files you closed. Choose the number that corresponds with the parameter file you want to open.

Exit command (File menu)

Use this command to end your PC1D session. You can also use the Close command on the application Control menu. PC1D prompts you to save parameter files with unsaved changes.

Shortcuts

Mouse: Double-click the application's Control menu button.
Keys: ALT+F4

Toolbar command (View menu)

Use this command to display and hide the Toolbar, which includes buttons for some of the most common commands in PC1D, such as File Open. A check mark appears next to the menu item when the Toolbar is displayed.

See [Toolbar](#) for help on using the toolbar.

Toolbar



The toolbar is displayed across the top of the application window, below the menu bar. The toolbar provides quick mouse access to many tools used in PC1D.

To hide or display the Toolbar, choose Toolbar from the View menu (ALT, V, T).

For help on the use of a particular toolbar button, select the button at the far right edge of the toolbar, displaying a question mark and an arrow. The cursor will change to match this button. Point the new cursor to the button in question and press the mouse key. A Help file will open on-screen with information about the purpose of that button.

Status Bar command (View menu)

Use this command to display and hide the Status Bar, which describes the action to be executed by the selected menu item or depressed toolbar button, and keyboard latch state. A check mark appears next to the menu item when the Status Bar is displayed.

See Status Bar for help on using the status bar.

Status Bar

The status bar is displayed at the bottom of the PC1D window. To display or hide the status bar, use the Status Bar command in the View menu.

The left area of the status bar describes actions of menu items as you use the arrow keys to navigate through menus. This area similarly shows messages that describe the actions of toolbar buttons as you depress them, before releasing them. If after viewing the description of the toolbar button command you wish not to execute the command, then release the mouse button while the pointer is off the toolbar button.

The right areas of the status bar indicate the following:

Indicator	Description
Iteration Counter	During a numerical solution, this frame shows how many iterations have been completed toward the current solution.
Convergence Indicator / Element Counter	During a numerical solution, this frame shows a bar that moves towards the right as the numerical solution converges to an answer. Otherwise, this frame indicates how many finite elements are defined for the device.
Solution Status	This field indicates the status of the current numerical solution. If blank, then the current problem has not yet been solved, or the parameters have been changed since the last solution. Valid solutions are indicated as Equil, Steady, or Trans for Equilibrium, Steady-State, or Transient results. Do not use displayed results unless the Status frame indicates a valid solution.
Time Step Number	During a transient solution, this frame shows for which time step number the displayed graphs apply.

Index command (Help menu)

Use this command to display the opening screen of Help. From the opening screen, you can jump to step-by-step instructions for using PC1D and various types of reference information.

Once you open Help, you can click the Contents button whenever you want to return to the opening screen.

Using Help command (Help menu)

Use this command for instructions about using Help.

About command (Help menu)

Use this command to display the copyright notice and version number of your copy of PC1D.

Context Help command



Use the Context Help command to obtain help on some portion of PC1D. When you choose the Toolbar's Context Help button, the mouse pointer will change to an arrow and question mark. Then click somewhere in the PC1D window, such as another Toolbar button. The Help topic will be shown for the item you clicked.

Shortcut

Keys: SHIFT+F1

Title Bar

The title bar is located along the top of a window. It contains the name of the application and parameter file.

To move the window, drag the title bar. Note: You can also move dialog boxes by dragging their title bars.

A title bar may contain the following elements:

- Application Control-menu button
- Maximize button
- Minimize button
- Name of the application
- Name of the parameter file
- Restore button

Scroll bars

Displayed at the right and bottom edges of the window. The scroll boxes inside the scroll bars indicate your vertical and horizontal location in the window. You can use the mouse to scroll to other parts of the window.

Size command (System menu)

Use this command to display a four-headed arrow so you can size the active window with the arrow keys.



After the pointer changes to the four-headed arrow:

1. Press one of the DIRECTION keys (left, right, up, or down arrow key) to move the pointer to the border you want to move.
2. Press a DIRECTION key to move the border.
3. Press ENTER when the window is the size you want.

Note: This command is unavailable if you maximize the window.

Shortcut

Mouse: Drag the size bars at the corners or edges of the window.

Move command (Control menu)

Use this command to display a four-headed arrow so you can move the active window or dialog box with the arrow keys.



Note: This command is unavailable if you maximize the window.

Shortcut

Keys: CTRL+F7

Minimize command (application Control menu)

Use this command to reduce the PC1D window to an icon.

Shortcut

Mouse: Click the minimize icon  on the title bar.
Keys: ALT+F9

Maximize command (System menu)

Use this command to enlarge the active window to fill the available space.

Shortcut

Mouse: Click the maximize icon  on the title bar; or double-click the title bar.
Keys: CTRL+F10 enlarges a window.

Close command (Control menus)

Use this command to close the active window or dialog box.

Double-clicking a Control-menu box is the same as choosing the Close command.

Shortcuts

Keys: CTRL+F4 closes the program window
 ALT+F4 closes the active window or dialog box

Restore command (Control menu)

Use this command to return the active window to its size and position before you chose the Maximize or Minimize command.

Switch to command (application Control menu)

Use this command to display a list of all open applications. Use this "Task List" to switch to or close an application on the list.

Shortcut

Keys: CTRL+ESC

Dialog Box Options

When you choose the Switch To command, you will be presented with a dialog box with the following options:

Task List

Select the application you want to switch to or close.

Switch To

Makes the selected application active.

End Task

Closes the selected application.

Cancel

Closes the Task List box.

Cascade

Arranges open applications so they overlap and you can see each title bar. This option does not affect applications reduced to icons.

Tile

Arranges open applications into windows that do not overlap. This option does not affect applications reduced to icons.

Arrange Icons

Arranges the icons of all minimized applications across the bottom of the screen.

Modifying the Parameter File

PC1D Parameter files are modified in one of two ways:

1. You can use the menu structure to select various aspects of the device or excitation and open dialog boxes which allow you to change the values, or
2. You can double-click on most lines in the Parameter View to open a dialog box which will allow you to enter a new value for that parameter.

If the set of parameters are modified in any way, PC1D will prompt you whether you want to save the modified parameters, before it will allow you to exit the program or overwrite the current parameters.

No Help Available

No help is available for this area of the window.

No Help Available

No help is available for this message box.

Compute Menu Commands

The Compute menu offers the following commands:

Run	Begins numerical solution of the problem as currently set up, starting with equilibrium and progressing to steady-state or transient conditions as specified by the current Excitation Mode. <i>A button on the toolbar (running person) provides a quick way to invoke Run.</i>
Stop	Stops the numerical computation in progress. The computation is allowed to continue for one second after this command is selected. If it has not then converged, the computation is terminated.
Continue	Resumes the numerical computation that was interrupted by the Stop command. Or, if the excitation for a problem has been changed since it converged, PC1D will redo the last solution phase using the current solution as a starting point. You can use this feature to help convergence in many cases.
Single Step	Computes only the next step in the solution. A “step” is defined here as the unit of computation leading to the next display of graphical information. If the Numerical dialog has been set to display plots after every iteration, then this is only one iteration. Otherwise, the computation proceeds until convergence is obtained for equilibrium, steady state, or one time step. <i>A button on the toolbar (stepping person) provides a quick way to Single Step.</i>
Start Again	Resets the computation so that a subsequent Single Step command will start with problem initialization and equilibrium solution. Note that the Run command always starts with problem initialization.
Batch	Opens a dialog box for enabling batch mode and identifying which tab-delimited ASCII file contains the batch-parameter information.
Numerical	Opens a dialog box for setting the parameters that control the numerical method. These affect the number of finite elements and the convergence of the solution.

Numerical Command (Compute menu)

This command opens a dialog box which allows you to set the following parameters which affect the numerical computation algorithm. These parameters do not alter the definition of either the device or the excitation, but they will influence the accuracy of the solution and the speed with which the solution converges to an answer. Improper settings of these parameters can cause the program to fail to converge even for simple problems, so the default values of each parameter are listed here for reference.

Element Size Factor This value determines the size of the finite elements used to partition the device. A smaller element size factor produces smaller elements, which improves accuracy but takes more time. The default value is 0.5. Values greater than 1.0 can produce elements so large that convergence problems may arise. Once the factor is small enough that all 500 available elements are used, making it smaller will have no further effect.

Normalized Error Limit This value determines when the solution is said to have converged. Iterations will continue until the largest change in any of the three potentials (electron or hole quasi-Fermi potential, or electrostatic potential) at any node is less than this factor times the thermal voltage, kT/q (which is about 26 mV at room temperature). The default value is $1E-6$, and will rarely need modification.

Normalized Potential Clamp This value determines the maximum change in one iteration that is allowed for any potential at any node, as a multiple of the thermal voltage, kT/q . The default value is 1, which is rather conservative, favoring robustness over speed. Increase this value for more speed if convergence is not a problem. In particular, a higher value may be desirable to increase speed when large reverse-bias voltages are imposed. Note that you can change the clamp value during a solution, for those times when you need a small value for steady state but a larger value will do for subsequent transient steps.

Clamping Phi/Psi These check boxes determine how the Normalized Potential Clamp is imposed. Selecting Psi clamping prevents the electrostatic potential from changing by more than the clamp amount with each iteration. Selecting Phi prevents the separation between each quasi-Fermi potential and the electrostatic potential from changing too much. **AT LEAST ONE OF THESE BOXES SHOULD BE SELECTED**, as many problems will “blow up” very quickly without some form of clamping imposed. The default is for both boxes to be selected.

Maximum Time This value determines how long (in seconds) a solution is permitted to continue before it is judged to be non-convergent. The default value is 60 seconds. Some problems involving internal shunt elements may take longer than this to converge, but most problems that haven't converged after a minute never will.

Renode These three check boxes determine when the finite-element nodes are reallocated. When initialized, the problem has 100 elements in each defined region. It is usually desirable to renode during the equilibrium solution, since the initial node allocation does not yet know where the critical junction regions are located. Selecting renode for the steady-state solution serves to refine the previous node allocation based on the equilibrium solution, and allows the program to do an emergency renode if the quasi-Fermi potential step across any one element exceeds 32 times the thermal voltage. Renoding for a transient solution should only be invoked when the excitation significantly alters the space-charge regions, and it should especially be avoided during fast transients where time derivatives are important. The default is to renode during equilibrium and steady state solutions, but not during transient solutions.

Display graphs after every iteration This check box is provided for those who are interested in observing the numerical computation performed by PC1D. When checked, the plots on screen are updated after every iteration, not just when the solution has converged. This slows the solution considerably, so should be chosen only when the convergence behavior is of interest.

Total velocity
saturation

This check box is provided for those who need to impose rigid saturation of the total velocity of the carriers, due to both drift and diffusion. Normally, when this box is not selected, the mobility of the carriers is reduced only in response to a high electric field. When this box is selected, the mobility is reduced in response to a high gradient in the quasi-Fermi potential. Although checking this box invokes the more-correct physical limitation, convergence is more difficult and solutions can take more than twice as long to complete. It is recommended only for heterostructure devices in which the current is limited by thermionic emission over an energy barrier.

Options menu commands

The Options menu provides commands that affect the environment in which PC1D runs on your computer, but which have no effect whatever on the problem or its solution. The features you select with these commands are stored on your computer in a file named PC1D.INI which is normally stored in your WINDOWS directory. Consequently, they will be used whenever you start PC1D on your computer. They will not apply if you save a problem file on disk and subsequently work with it using a different computer.

File Locations	This command opens a dialog box which allows you to specify which subdirectory you would like to serve as the default location for each type of external file used by PC1D. By segregating your files into separate directories, you will find it much easier to locate the files that you need later. If a field is left blank, the default location is assumed to be the directory from which PC1D was started.
Store Node Data	This command, when ticked, instructs PC1D to store the current solution on disk whenever it stores the current problem parameters. This is useful if you are in the middle of examining a solution in detail but need to quit the program for a period. Saving the node data preserves the most recent solution details, preventing you from needing to repeat the solution when you return, but it also increases the size of the saved PRM files substantially. The default is for this option to be disabled.
Device Update	This command, when ticked, instructs PC1D to update all of the on-screen graphs whenever a change is made in any of the parameters that define the device. This can be quite helpful when setting up the device parameters, because you can immediately see the impact of a change in doping, thickness, etc. The default for this option is for it to be enabled. However, some computers may not be fast enough to keep up. If your computer seems to be lagging behind you when you are setting up a problem, consider disabling this option.

PC1D Physical Constants

The following physical constants are used in PC1D calculations:

Elementary charge, q 1.6021773E-19 C

Thermal voltage at 300 K, kT/q 0.025851483 V

Photon energy factor, hc 1239.8424 eV-nm

Permittivity of free space, ϵ_0 8.8541878E-14 F/cm

Kelvin-Celsius offset 273.15 K

Circle geometry constant, π 3.14159265359

The following program constants are constraints in this version of PC1D:

Maximum Elements 500

Maximum Time Steps 200

Maximum Wavelengths 200

Maximum Regions 5

Maximum Batch Columns 30

Device Schematic

A schematic diagram of the device is displayed in the parameter view, providing visual feedback of doping, texturing, and internal shunt elements. This makes it easier to recognize mistakes when designing complex devices. You can also double-click on parts of the diagram to change various parameters.

The parameters which can be modified from the diagram are:

Parameter	Where to click
Background Doping	Centre of a region
First Front Diffusion	Front 3mm of a region
First Rear Diffusion	Rear 3mm of a region
Front/Rear Texturing	Within 3mm outside the top/bottom of the cell
Surface Charge	Between 3 and 6mm of the top/bottom of the cell
Region Thickness	Within 3mm outside the right of the cell
Shunt elements	Within 3mm outside the left of the cell
Contacts	To the left of the shunt elements

Print command (File menu)

Use this command to print the current window. This command presents a Print dialog box, where you may specify the range of pages to be printed, the number of copies, the destination printer, and other printer setup options.

Shortcuts

Toolbar:



Keys:

CTRL+P

Print dialog box

The following options allow you to specify how the current window should be printed:

Printer

This is the active printer and printer connection. Choose the Setup option to change the printer and printer connection.

Setup

Displays a Print Setup dialog box, so you can select a printer and printer connection.

Print Range

Specify the pages you want to print:

All Prints the entire current window.

Selection Prints the currently selected text.

Pages Prints the range of pages you specify in the From and To boxes.

Copies

Specify the number of copies you want to print for the above page range.

Collate Copies

Prints copies in page number order, instead of separated multiple copies of each page.

Print Quality

Select the quality of the printing. Generally, lower quality printing takes less time to produce.

Print Progress Dialog

The Printing dialog box is shown during the time that PC1D is sending output to the printer. The page number indicates the progress of the printing.

To abort printing, choose Cancel.

Print Preview command (File menu)

Use this command to display the active current window as it would appear when printed. When you choose this command, the main window will be replaced with a print preview window in which one or two pages will be displayed in their printed format. The print preview toolbar offers you options to view either one or two pages at a time; move back and forth through the current window; zoom in and out of pages; and initiate a print job.

Print Preview toolbar

The print preview toolbar offers you the following options:

Print

Bring up the print dialog box, to start a print job.

Next Page

Preview the next printed page.

Prev Page

Preview the previous printed page.

One Page / Two Page

Preview one or two printed pages at a time.

Zoom In

Take a closer look at the printed page.

Zoom Out

Take a larger look at the printed page.

Close

Return from print preview to the editing window.

Print Setup command (File menu)

Use this command to select a printer and a printer connection. This command presents a Print Setup dialog box, where you specify the printer and its connection.

Print Setup dialog box

The following options allow you to select the destination printer and its connection.

Printer

Select the printer you want to use. Choose the Default Printer; or choose the Specific Printer option and select one of the current installed printers shown in the box. You install printers and configure ports using the Windows Control Panel.

Orientation

Choose Portrait or Landscape.

Paper Size

Select the size of paper that the current window is to be printed on.

Paper Source

Some printers offer multiple trays for different paper sources. Specify the tray here.

Options

Displays a dialog box where you can make additional choices about printing, specific to the type of printer you have selected.

Network...

Choose this button to connect to a network location, assigning it a new drive letter.

Page Setup command (File menu)

<< Write application-specific help here. >>

Batch Mode (Compute menu)

Batch mode is a short cut which allows you to rapidly perform an optimization study for a particular configuration (PRM file). Rather than creating a series of PRM files, you only need to create one, and then specify which parameters should be varied.

To do a batch run, click the 'Batch' button on the toolbar (the icon shows many people running).

Parameters can be chosen from the drop-down lists. There's a description of the parameter at the bottom of the dialog box. For input parameters, you need to specify the range to be varied over, the number of different values, and if they should be varied logarithmically or linearly.

When you've finished, press OK. The table will appear at the bottom of the parameter view. Results parameters will appear as question marks.

Run the simulation as normal. As each result is calculated, it will appear in the table, replacing the question mark. The number of simulations performed so far will be displayed in the title bar.

Examining the results

Use the Copy button to copy the results into the clipboard. From there, they can be pasted into other programs. For example, the results could be pasted into a spreadsheet and graphed.

If you're interested in how a particular graph varies, select it in Interactive Graph View before running the batch. After the batch has finished, use the PageUp and PageDown keys to see how the graphs varied for different parameter values. (You can keep the axes constant while doing this by selecting 'retain zoom' in the Graph menu. You can reset the graphs using 'reset history graphs' in the same menu).

Note: Any parameter that is disabled in the PRM file will have no effect on batch results. For example, the front texture angle is irrelevant if texturing is disabled.

Advanced feature: Permute

If you vary multiple parameters, they can be varied together, or you can solve for all combinations. Click the 'Permute' check box to do all permutations of that input parameter. If permute is off, the parameter will be tied to the parameter above it.

External Batch Files

If you find the 'QuickBatch' method is too restrictive, you can also use an external batch file. This file should contain only the parameters that vary between simulations, and the results that are desired. For example, a batch file for a solar cell could contain a list of background doping values, and Voc, Isc and Pmax as desired results. To perform a batch run, do the following steps:

From a spreadsheet (e.g. Microsoft Excel), enter the parameters in tabular form. Place input parameters on the left hand side of the table, and results parameters on the right. The parameters can be chosen from a list of about 150. (See below for the list).

There is no intrinsic limit to the length of a batch file. Previously, batch files were limited to 100 lines. The maximum number of lines now depends on the operating system you are using. In Windows 95, batch files are limited to a few hundred lines (more if smaller font sizes are used in Parameter View). This limitation doesn't apply to Windows NT.

Batch files can have 30 fields across (although only the first 12 will print on A4 paper, unless you select a very small font size or copy the results into a spreadsheet program).

Save the table in "Tabbed Text" format. The file should end in ".txt".

From PC1D, open your PRM file. Using the Compute:Batch menu, select the text file that you created with your spreadsheet.

Note: Shortcuts to files on network drives are not supported. You must select the batch file directly.

Example

There is a sample batch file included with PC1D called BatchExample.txt, which may be useful for first-time users of the batch capability.

Batch Parameters

The parameters which can be specified are:

Input Parameters

Device Parameters

Region Parameters

Excitation Parameters

Numerical Parameters

Results Parameters

Voltage, current and power

Spatial results

Several abbreviations are used in batch parameter titles:

Fr = Front, Rr = Rear

Tx=Texture

Refl = Reflection, BroadRef=Broadband reflection

Bkgnd=Background

Dop=Doping

Pos=Position

Pri=Primary, Sec=Secondary

Insy=Intensity

Mono=Monochrome wavelength

Coll=Collector

SS=Steady-state value, TR1=Initial transient value, TR2=final transient value

Input parameters

Device parameters

These batch parameters specify the device parameters which apply to the entire device (not just individual regions). This includes the front and rear surface charge and reflection properties, and the device area.

<i>Name</i>	<i>Meaning</i>	<i>Units</i>
Area	Area	cm ²
FrTxAngle	Front texture angle	degrees
FrTxDepth	Front texture depth	um
RrTxAngle	Rear texture angle	degrees
RrTxDepth	Rear texture depth	um
FrBarrier	Height of front surface barrier (positive= bands bend up)	eV
FrCharge	Front surface charge	cm ⁻²
RrBarrier	Height of rear surface barrier (positive=eV bands bend up)	eV
RrCharge	Rear surface charge	cm ⁻²
EmitterR	Emitter internal resistance	ohms
BaseR	Base internal resistance	ohms
CollectorR	Collector internal resistance	ohms
EmitterX	Emitter distance from front	um
BaseX	Base distance from front	um
CollectorX	Collector distance from front	um
Shunt1	Value of 1st shunt element	seimens
Shunt2	Value of 2nd shunt element	seimens
Shunt3	Value of 3rd shunt element	seimens
Shunt4	Value of 4th shunt element	seimens
Shunt1Xa	Anode position of 1st shunt element	um
Shunt1Xc	Cathode position of 1st shunt	um
Shunt2Xa	Anode position of 2nd shunt	um
Shunt2Xc	Cathode position of 2nd shunt	um
Shunt3Xa	Anode position of 3rd shunt	um
Shunt3Xc	Cathode position of 3rd shunt	um
Shunt4Xa	Anode position of 4th shunt	um
Shunt4Xc	Cathode position of 4th shunt	um
FrRefl	Front reflectance (fixed)	%
RrRefl	Rear reflectance (fixed)	%
FrBroadRef	Front broadband reflectance	%
RrBroadRef	Rear broadband reflectance	%
FrOutThick	Thickness of outer layer (for front broadband)	nm
FrMidThick	“ front middle layer	nm
FrInThick	“ front inner layer	nm
FrOutIndex	Refractive index of front outer layer	-
FrMidIndex	“ front middle layer	-

FrInIndex	“ front inner layer	-
RrOutThick	Thickness of rear outer layer	nm
RrMidThick	“ rear middle layer	nm
RrInThick	“ rear inner layer	nm
RrOutIndex	Refractive index of rear outer layer	-
RrMidIndex	“ rear middle layer	-
RrInIndex	“ rear inner layer	-
FrIntRefl1	Front:First internal reflection	%
FrIntRefl2	Subsequent internal reflection	%
RrIntRefl1	Rear: First internal reflection	%
RrIntRefl2	Subsequent internal reflection	%
FrIntRefl	Front internal reflection (all passes)	%
RrIntRefl	Rear internal reflection (all passes)	%

Region parameters

These batch parameters let you define region-specific device parameters.

With these parameters, you must specify a region number (in parentheses) after the name. e.g. BkgndDop(3) gives values for the background doping of region 3. Region 1 is the region closest to the front of the device.

<i>Name</i>	<i>Meaning</i>	<i>Units</i>
Thickness	Thickness of region	um
BkgndDop	Background doping	cm-3
FrDopPeak1	1st Front doping - peak value	cm-3
FrDopDpth1	“ - depth factor	um
FrDopPos1	“ - peak position	um
FrDopPeak2	2nd front diffusion - peak	cm-3
FrDopDpth2	“ - depth factor	um
FrDopPos2	“ - peak position	um
RrDopPeak1	1st rear diffusion -peak	cm-3
RrDopDpth1	“ - depth factor	um
RrDopPos1	“ - peak position	um
RrDopPeak2	2nd rear diffusion - peak	cm-3
RrDopDpth2	“ - depth factor	um
RrDopPos2	“ - peak position	um
BulkTau	Bulk recomb.: electron lifetime	us
BulkTaup	Bulk recombination: hole lifetime	us
BulkTau	Bulk recombination, set electron and hole lifetime to the same value	us
BulkEt	Bulk trap energy level	eV
FrSn	Front surface electron recombination velocity	cm/s
FrSp	Front surface hole recombination velocity	cm/s
FrEt	Front surface trap energy level	eV
RrSn	Rear surface electron recombination velocity	cm/s
RrSp	Rear surface hole recombination velocity	cm/s

RrEt	Rear surface trap energy level	eV
FrS	Front surface recomb velocity, electrons and holes	cm/s
RrS	Rear surface recomb velocity, electrons and holes	cm/s
Ni200	Intrinsic concentration at 200K	
Ni300	Intrinsic concentration at 300K	
Ni400	Intrinsic concentration at 400K	
BandGap	BandGap	eV
AbsEd1	1st Direct BandGap (for absorption)	eV
AbsEd2	2nd Direct Bandgap	eV
AbsEi1	1st Indirect BandGap	eV
AbsEi2	2nd Indirect BandGap	eV

Excitation parameters

These parameters give you control over the excitation parameters to be used in each individual run in a batch mode. These are the same as the parameters in the dialog boxes which can be accessed from the Excitation menu.

<i>Name</i>	<i>Meaning</i>	<i>Units</i>
Temp	Temperature of device	kelvin
BaseResSS	Base steady-state resistance	ohms
BaseResTR	Base transient resistance	ohms
BaseVltSS	Base steady-state voltage	V
BaseVltTR1	Base transient initial voltage	V
BaseVltTR2	Base transient final voltage	V
CollResSS	Collector steady-state resistance	ohms
CollResTR	Collector transient resistance	ohms
CollVltSS	Collector steady-state voltage	V
CollVltTR1	Collector transient initial voltage	V
CollVltTR2	Collector transient final voltage	V
PriInsySS	Primary source steady-state intensity	Wcm ⁻²
PriInsyTR1	Primary source initial transient intensity	Wcm ⁻²
PriInsyTR2	Primary source final transient intensity	Wcm ⁻²
PriMonoSS	Pri: Wavelength - steady state	nm
PriMonoTR1	Pri: initial transient wavelength	nm
PriMonoTR2	Pri: final transient wavelength	nm
PriBlackT	Pri: Blackbody temperature	kelvin
SecInsySS	Secondary source steady-state intensity	Wcm ⁻²
SecInsyTR1	Sec: Initial transient intensity	Wcm ⁻²
SecInsyTR2	Sec: Final transient intensity	Wcm ⁻²
SecMonoSS	Sec: Wavelength - steady state	nm
SecMonoTR1	Sec: Initial transient wavelength	nm
SecMonoTR2	Sec: Final transient wavelength	nm
SecBlackT	Sec: Blackbody temperature	kelvin

Numerical parameters

These parameters give you control over the numerical parameters to be used in each individual run in a batch mode.

These are the same as the parameters in the Compute:Numerical... dialog box. They're provided here to make it possible to do batch runs involving simulations with different convergence properties.

<i>Name</i>	<i>Meaning</i>	<i>Units</i>
ElemSize	Element size factor	-
ErrorLimit	Normalized error limit	-
Clamp	Normalized potential clamp	-

Results parameters

Voltage, current and power

These results parameters give you access to the voltage, current and power for the base and collector contacts, and for the internal shunt elements. The values displayed in the batch results will be the steady-state values (for steady state excitation mode) or the final transient time step (for transient excitation).

<i>Name</i>	<i>Meaning</i>	<i>Units</i>
Vb	Base voltage	V
Vc	Collector voltage	V
Ib	Base current	A
Vb	Collector current	A
BaseVoc	Voc, base contact	V
BaseIsc	Isc, base contact	A
BasePmax	Pmax, base contact	W
CollVoc	Voc, collector contact	V
CollIsc	Isc, collector contact	A
CollPmax	Pmax, collector contact	W
V1	Voltage across 1st shunt element	V
I1	Current through 1st shunt element	A
V2	Voltage across 2nd shunt element	V
I2	Current through 2nd shunt element	A
V3	Voltage across 3rd shunt element	V
I3	Current through 3rd shunt element	A
V4	Voltage across 4th shunt element	V
I4	Current through 4th shunt element	A

Spatial results

These parameters give you access to all of the results used by the spatial graphs . You must specify the distance from front (in um) in parentheses after the name. e.g. Jn(10.5) gives the electron current density at the point 10.5 um from the front surface. The values displayed in the batch results will be the steady-state values (for steady state excitation mode) or the final transient time step (for transient excitation).

<i>Name</i>	<i>Meaning</i>	<i>Units</i>
Na	Acceptor Doping Density	cm ⁻³
RBulk	Bulk Recombination Rate	cm ⁻³ /s
Rho	Charge Density	C/cm ³
Ec	Conduction Band Edge	eV
Cond	Conductivity	S/cm
CCum	Cumulative Conductivity	S
CCum_	Cumulative Excess Conductivity	S
Gcum	Cumulative Photogeneration	s ⁻¹
Rcum	Cumulative Recombination	s ⁻¹
Perm	Dielectric Constant (Permeability)	-
Ld	Diffusion Length	m
Nd	Donor Doping Density	cm ⁻³
Eg	Effective Energy Gap	eV
Nie	Effective Intrinsic Concentration	cm ⁻³

E	Electric Field	V/cm
In	Electron Current	A
Jn	Electron Current Density	A/cm ²
N	Electron Density	cm ⁻³
Ndiff	Electron Diff. Current Density	A/cm ²
Ndrift	Electron Drift Current Density	A/cm ²
MuN	Electron Mobility	cm ² /Vs
PhiN	Electron Quasi-Fermi Energy	eV
Vn	Electron Velocity	cm/s
Psi	Electrostatic Potential	V
Rho_	Excess Charge Density	C/cm ³
Cond_	Excess Conductivity	S/cm
N_	Excess Electron Density	cm ⁻³
Nratio	Excess Electron Density Ratio	
Psi_	Excess Electrostatic Potential	V
P_	Excess Hole Density	cm ⁻³
Pratio	Excess Hole Density Ratio	
PNratio	Excess pn Product Ratio	
G	Generation Rate	cm ⁻³ /s
Ip	Hole Current	A
Jp	Hole Current Density	A/cm ²
P	Hole Density	cm ⁻³
Pdiff	Hole Diff. Current Density	A/cm ²
Pdrift	Hole Drift Current Density	A/cm ²
MuP	Hole Mobility	cm ² /Vs
PhiP	Hole Quasi-Fermi Energy	eV
Vp	Hole Velocity	cm/s
Tau	LLI Carrier Lifetime	s
PNnorm	Normalized Excess pn Product	
Res	Resistivity	ohm*cm
It	Total Current	A
Jt	Total Current Density	A/cm ²
Evac	Vacuum Energy	eV
Ev	Valence Band Edge	eV

Excitation Menu

The excitation menu offers the following commands:

Mode	This command opens a dialog box which allows you to control whether excitation is applied, and if so, whether it is steady-state or transient.
Temperature	This command opens a dialog box which allows you to set the temperature of the device. You can specify the temperature in either kelvin or degrees Celsius. This temperature is also used to compute adjustments to the carrier mobilities, bulk and surface recombination, optical absorption, and the exponential factor for any internal shunt diodes. However, the saturation current density, conductance, and capacitance of internal shunt elements are NOT adjusted for temperature; you must adjust these values separately for each temperature.
Base/Collector Source Circuits	These commands open dialog boxes which allow you to specify the Thevinin-equivalent circuits for the base and collector. These circuits are only active if base and/or collector contacts have been Enabled for this device.
Photogeneration	This selection allows you to introduce photogeneration in the device. Illumination can be provided from a primary and/or secondary light source, or the photogeneration profile can be supplied from an external file.
Excitation Files	This set of commands (New, Open, SaveAs) allows you to create, retrieve, or store binary files that contain all of the parameters needed to define the excitation. These excitation files, with suffix EXC, can then be used with a variety of different devices. Since version 5.0, the graph definitions are also saved in the excitation files. <i>A toolbar button is provided to open the dialog box for retrieving a previously defined excitation file.</i>

Excitation Mode (Excitation menu)

The excitation mode line can be set to Equilibrium, Steady State, or Transient. The choice determines the final state the solution will achieve before stopping. The solution will always proceed by solving first for Equilibrium, then Steady State, then Transient.

If the 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 time steps. 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 initial time step, which immediately precedes $t=0$, can be set to a different value from the remaining steps. Setting this value small permits the simulation of an abrupt change in excitation at time $t=0$. The entry for Number of Time Steps is the number of time steps of duration Time Step Size that are desired. The total duration of the transient will be the product of these two values.

Base/Collector Source Circuit (Excitation menu)

There are two identical Thevinin-equivalent source circuits; one associated with the base contact and one associated with the collector contact. Both circuits contain a voltage source and a series resistance. Different values of voltage and resistance can be specified for steady state versus transient conditions. Changing the voltage between its steady-state and initial transient value causes a step change in voltage at $t=0$. Setting the final transient value different from the initial transient value causes the voltage to sweep linearly from the initial to the final value during the course of the transient solution.

The source series resistance affects 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 value of series resistance can be specified either in ohms, or in ohm-cm. In the latter case, the series resistance is adjusted for each solution based on the area of the device being simulated.

A constant-voltage condition at a contact is achieved by setting the corresponding source series resistance 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 (1 megohm is usually plenty).

When a simple linear voltage ramp is not adequate, time-dependent source-voltage values can be supplied from an external ASCII data file having a filename with suffix VLT. Each line in the file should contain three numerical values, separated by one or spaces or by a tab. The first value on each line should be the time (in seconds), the second value is the source voltage (in volts), and the third value should be the series resistance (in ohms). The time values must be monotonically increasing. The maximum number of lines read from the file is 200. Note that the time values in this file do not affect the time steps used in the solution, which are determined by the Mode command. Rather, voltage values for each time step are interpolated from this file.

Photogeneration (Excitation menu)

The photogeneration options allow you to apply photoexcitation to the device, either as a photogeneration profile supplied via an external file, or using one or both of two light sources which can illuminate either the front or rear surface of the device.

When using the light sources, PC1D internally calculates the photogeneration rate within the device. At each incident wavelength, after accounting for incident-surface reflection, the light is absorbed in the device using the absorption coefficients for each region. If enabled, some light will be lost due to free carrier absorption. If the device has texture, then the photons do not travel parallel to the solution direction (x). The direction they travel 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. If the internal reflectance for the surface opposite the incident surface is non-zero, then some photons reflect from that surface with either the same angle at which they arrived (specular) or randomly-directed (diffuse). If the internal reflectance at the incident surface is non-zero, then some of this reflected light gets trapped within the device until it eventually is either absorbed in the device or fails to be reflected from one surface of the other.

Photogeneration Profile

One photogeneration option is to supply an external ASCII file with suffix GEN that contains photogeneration information. This file should have two values on each line, separated by one or more spaces or by a tab. The first value is a position representing the distance of that location from the front surface, in μm . The second value is the cumulative photogeneration rate in the device between the front surface and that position, in carrier-pairs per second per square cm of projected area. 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. Photogeneration information provided for positions beyond the rear surface of the device are ignored. If the device is thicker than the last entry in the file, then no photogeneration is assumed beyond the last position defined in the file.

Primary/Secondary Illumination Intensity

These commands open a dialog box to examine or modify the magnitude and time dependence of light incident on the device, and to select whether the light is incident on the front or rear surface of the device. To use either of the illumination sources, you must first Enable that source from within this dialog box. Different values of illumination source intensity can be specified for steady state versus transient conditions. Changing the intensity between its steady-state and initial transient value causes a step change in intensity at $t=0$. Setting the final transient value different from the initial transient value causes the intensity to sweep linearly from the initial to the final value during the course of the transient solution. The values entered correspond to the total power density normally incident on the 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 simple linear ramp of intensity is inadequate, time-dependent light intensity values can be entered from an external ASCII data file having a filename suffix LGT. Each line in the file should contain two values, separated by one or more spaces or by a tab. The first is a time, in seconds. The second is an intensity value with units of W/cm^2 . The lines must have monotonically increasing time values. The maximum number of lines in the file is 200. The time values in this file do not affect the time step size or number of time steps used in the solution. Rather, an intensity value is interpolated from the LGT file for each point in time specified in the Mode command.

Primary/Secondary Illumination Spectrum

The illumination spectrum choices are Monochrome, Black-Body, or External.

The Monochrome option allows you to specify that all of the incident power occurs at a single wavelength. Different values of wavelength can be specified for steady state versus transient conditions. Changing the wavelength between its steady-state and initial transient value causes a step change in wavelength at $t=0$. Setting the final transient value different from the initial transient value causes the wavelength to sweep linearly from the initial to the final value during the course of the transient solution.

The Black-Body option allows you to invoke a black-body spectrum corresponding to a specified temperature. The spectrum is actually implemented as a group of discrete wavelengths, so you must declare the number of discrete wavelengths (maximum 200) and the range of wavelengths to include. You can artificially limit the wavelength range to simulate the effect of a filtered spectrum. The total power density in the black-body spectrum incident on the device is adjusted to match the values specified for illumination intensity. Note that for a given temperature, there is a limit to the intensity that can be obtained from a black body source; however, PC1D does not verify whether this limit has been exceeded.

The External option allows you to supply an external ASCII file that defines a spectral distribution, represented as a group of discrete wavelengths. These files have suffix SPC. Several important spectrums are provided with PC1D, including the air-mass 1.5 direct and global ASTM solar spectrums, and the extraterrestrial solar spectrum. These files contain two values on each line, separated by one or more spaces or a tab. The first value is a wavelength in nm, and

the second is a power density in W/cm² (NOT spectral density, W/μm/cm²). The entries must be in order of increasing wavelength, with a maximum of 200 wavelengths allowed. The power densities are scaled as necessary so that the total light intensity for the spectrum as a whole equals the value specified for illumination intensity.

Device Menu

The Device menu offers selections in four groups:

Region Manager	This group of selections allow you to manipulate the number of regions in your device. Each region can be a different material, with its own parameters.
Region Parameters	This group of selections pertain only to the current region. They allow you to examine and modify the thickness, material, doping, and recombination parameters.
Device Parameters	This group of selections pertains to the device as a whole. They allow you to examine and modify the area of the device, surface texture, surface charge, external circuit contact locations, and the optical reflectance, both external and internal.
Device Files	This group of commands (New, Open, SaveAs) allows you to create, retrieve, and store device files, having a suffix DEV. These contain all of the information about the device only, with no excitation or numerical solution data.

Region Manager (Device menu)

The Region Manager selections in the Device Menu offer the following commands:

- Select Region Selecting a region number makes that the “current” region, such that any examination or modification of region parameters using the Device Menu will apply to that region. A region can also be selected by clicking on any line pertaining to that region in the Parameter View. *The current region is displayed on the toolbar as a button. Pressing this button increments the current region.* Region numbers that are not defined for this device are disabled.

- Insert Region Selecting a number inserts a new region at that point in the device. Region 1 is always the front of the device. The new region will be defined with parameters that duplicate the most recently selected region, except that surface recombination and diffusions are set to zero. Once inserted, the new region becomes the current region.

- Remove Region Selecting a number removes that region from the device. Any higher-numbered regions are shifted toward the front of the device. The current region is reset to one of the regions still remaining.

Region Parameters (Device menu)

The Region Parameter selections in the Device menu allow you to examine or modify the parameters in the following categories. In each case, the command applies only to the region that is currently selected. *The currently selected region is indicated on the toolbar as a button.*

Thickness

Material

Doping

Recombination

Region Thickness (Device menu)

This command opens a dialog box for specifying the thickness of the current region, in units of μm .

Region Material (Device menu)

The parameters that define a type of material can be stored in binary files having the suffix MAT. Three commands at the bottom of the Material menu allow you to create, retrieve, and store these material files. Several MAT files for the most commonly encountered materials are provided on the distribution diskette with PC1D. Loading a new MAT file will only affect the current region. In addition, the Material menu allows you to examine and modify the following parameters for the current region:

Permittivity

This is the low-frequency (dc) value of the permittivity, expressed as a multiple of the permittivity of free space. This value is also known as the dielectric constant of the material. It defines the amount of charge needed to produce a given electric field in the material.

Band Structure

This command opens a dialog box which contains several parameters that are intrinsic to the band structure of the material.

The Band Gap value is used with the Electron Affinity to establish the location of the conduction and valence band edges relative to the vacuum level. For single-region problems they have no effect on the numerical solution, but for heterostructures they affect the discontinuity in the band edges at the interfaces between regions in accordance with an Anderson model of the interface, producing charge dipoles at these locations even if the doping is uniform.

The Nc/Nv Ratio refers to the effective density of states for the conduction and valence bands. This ratio affects the location of the intrinsic energy level within the bandgap, and only affects the solution of heterostructure problems.

Three values can be entered for the intrinsic carrier concentration, at 200, 300 and 400 K. These values are used to interpolate a value for the solution temperature, using a parabolic fit of $\log(n_i)$ versus $1/T$. These values have a major impact on the results obtained for any minority-carrier device.

Material Recombination

This dialog box allows you to specify parameters that affect how the recombination rates in this material are influenced by temperature, doping, and electric field. These parameters are often fairly consistent for a given material, and thus are included as material parameters. The base recombination rates for surfaces and bulk are not considered to be a property of the material, and are set separately for each region. See the entry for Region Recombination.

Three Auger recombination coefficients are required to properly compute recombination in either heavily doped or highly injected material. The n-type and p-type coefficients apply to material in low-level injection, while the high-injection coefficient applies to material in high-level injection. The high-injection coefficient represents the sum of both n-type and p-type Auger recombination. It is included as a single factor because in high-level injection the excess electron and hole concentrations are similar. The recombination rate due to Auger processes is computed from the following expression, which weights the result to favor either the low-injection or high-injection coefficients as appropriate:

$$R = (C_n n + C_p p)(pn - n_{ie}^2),$$
$$C_n = C_{nLLI} \left(\frac{N_D}{N_D + p} \right) + \frac{C_{HLI}}{2} \left(\frac{p}{N_D + p} \right),$$
$$C_p = C_{pLLI} \left(\frac{N_A}{N_A + n} \right) + \frac{C_{HLI}}{2} \left(\frac{n}{N_A + n} \right).$$

A band-to-band recombination coefficient can be provided. This is primarily important only for direct-bandgap materials. The additional recombination at each point due to coefficient B is given by

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

Parameters can be entered for the material to indicate the effect of temperature and doping on SRH surface and bulk recombination. For bulk lifetime, the values of τ_n and τ_p specified for the region are modified for use in the SRH recombination formula to reflect the actual device temperature and the local doping density. These parameters can be adjusted to either increase or decrease the impact of temperature and doping. The temperature effect is given by the following formula, applied to both electron and hole lifetimes:

$$\tau(T) = \tau_o \left(\frac{T}{300K} \right)^\alpha,$$

and the doping effect is given by

$$\tau(N) = \tau_o \left(\frac{N}{N_{onset}} \right)^\alpha; \quad N > N_{onset}$$

The parameters in these two expressions are empirical and very dependent on the crystal growth technology. For silicon, reasonable choices are -0.5 for the temperature exponent, -0.4 for the doping exponent, and 1E15 cm⁻³ for the onset.

For surface recombination velocities, both electron and hole values are modified using the following formulas, which have the same form as the bulk-lifetime expressions.

$$S(T) = S_o \left(\frac{T}{300K} \right)^\alpha,$$

and the surface doping effect is given by

$$S(N) = S_o \left(\frac{N}{N_{onset}} \right)^\alpha; \quad N > N_{onset}$$

The parameters in these two expressions are very dependent on the surface treatment process. For thermal oxide on n-type silicon, reasonable choices are 0.5 for the temperature exponent, 1.0 for the doping exponent, and 1E18 cm⁻³ for the onset. No consistent dependence of S on surface doping density has been demonstrated for p-type silicon at the time of this writing.

Field-Enhanced Recombination

The Hurkx model for trap-assisted tunneling is used. This extends the conventional SRH recombination expression to include trap-assisted tunneling enhanced by the local electric field. This model replaces the carrier lifetimes τ_n and τ_p

$$\text{in the SRH model with } \frac{\tau_n}{1 + \Gamma}, \frac{\tau_p}{1 + \Gamma}$$

where Γ is defined by:

$$\Gamma = \text{prefactor} \cdot \frac{|F|}{F_\Gamma} \exp \left[\left(\frac{F}{F_\Gamma} \right)^2 \right] \quad (\text{equation A2 in Hurkx et. al.})$$

This expression can be parametrised by specifying a prefactor ($\text{prefactor} = 2\sqrt{3\pi} \cong 6.14$) and the value of F_Γ at 300K. Hurkx gives:

$$F_\Gamma = \frac{\sqrt{24m^*(kT)^3}}{q\hbar} \quad (\text{equation A3 in Hurkx et. al.})$$

and suggests $m^*=0.25m_0$.

WARNING: Do not use this model blindly!

This model is only valid provided the electric field does not get too large. It will be valid in general for midgap traps, but for shallow traps in regions of very high electric field, it may dramatically overestimate the impact of trap-assisted tunneling (for high fields, Γ no longer depends exponentially on F). In these situations, the value for the prefactor and for F_Γ can be modified to give more realistic simulations.

At present, there is insufficient experimental data to justify a more complex model, although Hurkx provides an additional high-field equation (eqn A5 of his 1992 paper).

Reference:

G.A.M. Hurkx, D.B.M. Klaassen, and M.P.G. Knuvers, "A New Recombination Model for Device Simulation Including Tunneling", IEEE Trans. Electron Devices, ED-39, (Feb. 1992), pp. 331-338.

Bandgap Narrowing

The parameters entered for bandgap narrowing allow the equilibrium pn product to depend on the local doping density according to the following expression:

$$pn = n_{ie}^2 = n_i^2 e^{(\Delta E_c + \Delta E_v)/kT},$$

$$\Delta E_c = Slope_N \times \ln\left(\frac{N_D}{Onset_N}\right),$$

$$\Delta E_v = Slope_P \times \ln\left(\frac{N_A}{Onset_P}\right),$$

where the log expressions are taken equal to zero for doping densities less than the respective onset values.

Mobilities

The mobilities of electrons and holes can either be set to fixed values throughout the region, or calculated based on the temperature and local doping density using a model. The model uses a different set of coefficients where the carrier is in the minority than where it is in the majority. The expression used has the following form:

$$\mu(x, T) = \mu_{\min} T_n^{\beta 1} + \frac{(\mu_{\max} - \mu_{\min}) T_n^{\beta 2}}{1 + \left(\frac{N_D(x) + N_A(x)}{N_{ref} T_n^{\beta 3}}\right)^{\alpha T_n^{\beta 4}}},$$

where T_n is the temperature normalized to 300 K. The mobility model also allows for a maximum velocity for each carrier. This is implemented as a reduction in the mobility at each point in the device where the velocity that the carriers would have achieved approaches or exceeds the saturation velocity. The low-field mobility is replaced with a high-field mobility given by:

$$\mu_{HF} = \frac{\mu_{LF}}{\sqrt{1 + \left(\frac{\mu_{LF} E}{v_{sat}}\right)^2}}.$$

Normally, the electric field is used for "E". This produces the required mobility reduction in regions where the high carrier velocity is due to a high electric field, as at junctions. A more accurate approach is to use the gradient in the quasi-Fermi potential in place of E in this expression. This can be done in PC1D by selecting *Total velocity saturation* in the Compute:Numerical dialog box. Note, however, that convergence is less robust and solutions will take much longer with this option selected.

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

Refractive Index

The index of refraction can be either a fixed value or a function of wavelength using data imported from an external disk file. External index-of-refraction files are standard ASCII files with a filename suffix INR. Each line in the file should contain two numerical values, separated by one or more spaces or a tab. The first value on each line is a value of free-space wavelength, in nm. The second value on each line is the index of refraction at that wavelength. The maximum number of lines allowed is 200. 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.

Optical Absorption Coefficients

The optical absorption coefficient as a function of wavelength can be generated as a function of temperature using an internal model, or read from an external disk file. When experimental data for the temperature of interest is available, it is generally better to use this data than to rely on the internal model. Only absorption that leads to photogenerated carriers should be considered, since PC1D assumes that every intrinsically absorbed photon produces an electron-hole pair. Parasitic absorption processes can be modelled through the Free-Carrier Absorption dialog.

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. The entries required in the absorption dialog are the energy gaps at 300 K. All four energy gaps are assumed to depend on the absolute temperature (T) through a temperature coefficient (β) and temperature offset (γ):

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

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

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

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

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

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 together 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 suffix ABS. Each line in the file should contain two numerical values, separated by one or more spaces or a tab. The first value on each line is a value of free-space wavelength, in nm. The second value on each line is the absorption coefficient, in cm^{-1} . The maximum number of lines allowed is 200. 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, PC1D will assume that the final (nonzero) value applies even out to very long wavelengths, thereby greatly overestimating the total photogeneration. Room-temperature absorption files are provided with PC1D for several materials.

Free-Carrier Absorption

Free-carrier absorption competes with intrinsic absorption, reducing the fraction of photons which generate electron-hole pairs.

The free-carrier absorption across a device is not only dependent upon the wavelength of incident light, but also depends upon the profile of the carrier concentrations across the device. It can be specified by the free-carrier absorption coefficient, which for monochromatic incident light is assumed to take the form

$$\alpha_{FC} = K_1 n \lambda^a + K_2 p \lambda^b$$

where K_1 , K_2 , a , b are empirically determined constants, and λ is measured in nm. At the time of writing, estimates of these parameters were available for several materials:

	K1	a	K2	b	Source
AlSb	1.9e-24	2			Fan
GaAs	4e-29	3			Fan

GaP	(1.5e-24)	(1.8)			Fan
GaSb	9e-31	3.5			Fan
Ge	~5e-25	~2			Fan
InAs	6.5e-29	3			Fan
InP	5e-27	2.5			Fan
InSb	2.8e-25	2			Fan
Si	2.6e-27	3	2.7e-24	2	Schmid

References:

D.K. Schroder, R.N. Thomas, and J.C. Swartz, "Free Carrier Absorption in Silicon", IEEE Trans. Electron Devices, ED-25, (Feb. 1978), pp. 254-261.

P.E. Schmid, "Optical Absorption in Heavily Doped Silicon", Phys. Review B (1981), Vol 23, pp. 5531-5536.

H.Y. Fan, "Semiconductors and Semimetals", ed. R.K. Williamson and A.C. Beer, Academic Press (1967), Vol 3, p.409.

For reasons of speed, PC1D calculates the photogeneration only once for each solution step, using the values of n and p which were calculated for the previous step. For example, the steady-state photogeneration will be based on the equilibrium concentrations of electrons and holes. Although n and p will vary from their equilibrium values, the perturbation is unlikely to have a significant effect.

For the rare situations where the change in n and p during solution is important, the Compute:Continue menu command can be used at the conclusion of a steady-state problem. This will cause the photogeneration to be recalculated, using the values of n and p obtained from the steady-state solution. The steady state solution will then be resolved using the new photogeneration.

Region Doping (Device menu)

This selection allows you to specify the doping densities for donors and acceptors in the current region.

Background Doping

This command opens a dialog box where you can specify the doping density and the doping type (n-type or p-type) for background doping that is present uniformly throughout the layer. The mobilities at 300 K for both electrons and holes at this doping density are displayed for reference, as are the diffusivities. The mobilities are either fixed or variable with doping, as selected for this region's material). The majority-carrier mobility shown is used to calculate the resistivity. If you want the doping density that corresponds to some desired 300-K resistivity, then you can just type the desired resistivity directly into this field, then hit the Tab key to set the doping density accordingly.

Front/Rear Diffusions

Doping profiles can be specified for both the front and rear surfaces of the current region. These profiles can be generated using internal models or can be imported from an external disk file.

The internal model allows for two diffusions at each surface. After selecting either the first or second diffusion for the surface, a dialog box is opened for specifying the diffusion details. *The dialog box for the first diffusion on the front and rear surfaces can also be opened using a button on the toolbar.* Each diffusion must be Enabled within its dialog box before it becomes active. The toolbar button will appear to be depressed if that diffusion has been Enabled. Diffusions can be specified as n-type or p-type. Each diffusion can take one of four shapes, defined as follows:

Uniform
$$N(x) = N_o; \quad x_p < x < x_p + x_d$$

Exponential
$$N(x) = N_o e^{-|x-x_p|/x_d}$$

Gaussian
$$N(x) = N_o e^{-(x-x_p)^2/x_d^2}$$

Erfc
$$N(x) = N_o \operatorname{erfc}\left((x - x_p) / x_d\right)$$

where N_o is the Peak Doping, x_d is the Depth Factor, and x_p is the Peak Position, all of which are adjustable. Note that $N(x) = N_o$ for $x < x_p$ for the Erfc shape.

Also shown in the dialog box are the calculated sheet resistance (ohms/square) and junction depth. The junction depth is based on the currently specified background doping density for the region. The sheet resistance is calculated based on the diffusion profile up to this junction depth, using a temperature of 300 K and the majority-carrier mobility (either fixed or variable, as specified for this region's material). The calculation of the sheet resistance ignores conductance due to dopants from the background or from other diffusions.

To obtain a diffusion with a given junction depth and sheet resistance, you can type the desired sheet resistance in the space provided, then press the Tab key. The Peak Doping is adjusted to give the requested sheet resistance. If the junction depth does not match the desired value, type the desired value in that field, then press Tab to adjust the Depth Factor to obtain the desired junction depth. If the sheet resistance is affected, you can re-enter the desired value in that space, and repeat the cycle until convergence is obtained. Rarely are more than a few iterations required to get quite close to the desired combination.

External diffusion profiles are standard ASCII files with a filename suffix DOP. These can be created using a text editor or generated by another computer program. Each line in the file should contain three numeric values, each separated by one or more spaces or a tab. The first value on each line is a value of distance inward from the surface of the region, with units of μm . The second value on each line is the donor doping density at that position, with units of cm^{-3} . The third value on each line is the acceptor doping density, with units of cm^{-3} . The position values must start at 0 and increase monotonically. The maximum number of positions that may be defined in the file is 500.

Region Recombination (Device menu)

This selection allows you to specify the recombination parameters in the bulk and at both the front and rear surface of the region.

Bulk Recombination

The bulk recombination model used by PC1D supplies a single Shockley-Read-Hall (SRH) energy level with separate electron and hole lifetimes and an adjustable trap energy level within the bandgap. *The dialog box for specifying these parameters for the current region can also be opened using a button on the toolbar.* This toolbar button will appear to be depressed if the electron or hole lifetime has been altered from its default value of 1 ms.

$$R_{SRH} = \frac{pn - n_{ie}^2}{\tau_n \left(p + n_{ie} e^{-E_t/kT} \right) + \tau_p \left(n + n_{ie} e^{E_t/kT} \right)}$$

The trap level, E_t , is specified with respect to the intrinsic level, which lies near the middle of the bandgap. The electron and hole lifetimes specified at the top of the dialog box are the values corresponding to intrinsic material at a temperature of 300 K.

Also listed in the bulk-recombination dialog box is the excess-carrier lifetime in low-level injection and the corresponding diffusion length, for material with the background doping and reference temperature specified in the dialog box. These calculated values include the effect of band-to-band and Auger recombination, the coefficients for which are specified for the material's band structure in this region, and also the effects of background doping on the SRH lifetimes as specified in the material recombination dialog. To obtain a specific low-level-injection lifetime or diffusion length, simply type the desired number in the space provided. PC1D will calculate the intrinsic electron and hole lifetimes needed. Note that the properties of the material (band-to-band, Auger recombination, etc.) place an upper limit on the achievable lifetime and diffusion length for a given background doping density.

Front/Rear Surface Recombination

The surface recombination model used by PC1D supplies a single Shockley-Read-Hall (SRH) energy level with separate electron and hole lifetimes and an adjustable trap energy level within the bandgap. *The dialog box for specifying these parameters for the front or rear surface of the current region can also be opened using buttons on the toolbar.* These toolbar buttons will appear to be depressed if there is recombination at that surface. When either S_n or S_p is nonzero, the SRH surface recombination rate is given by

$$R_S = \frac{S_n S_p (pn - n_{ie}^2)}{S_p \left(p + n_{ie} e^{-E_t/kT} \right) + S_n \left(n + n_{ie} e^{E_t/kT} \right)}$$

The trap level, E_t , is specified with respect to the intrinsic level, which lies near the middle of the bandgap. The electron and hole surface recombination velocities specified at the top of the dialog box are the values corresponding to intrinsic material at the reference temperature specified in the dialog box. The values of S_n and S_p entered are modified for use in the recombination formula to reflect the actual device temperature and the surface doping density. The extent of this modification depends on the parameters listed in the material recombination dialog box.

Also listed in the surface-recombination dialog box is the effective minority-carrier surface recombination velocity for material with the surface doping at a temperature of 300 K. To obtain a specific minority-carrier surface recombination velocity, simply type the desired number in the space provided. PC1D will calculate the intrinsic electron and hole surface recombination velocities needed.

When high-level injection occurs at the surface, you can choose to either continue with the Shockley-Read-Hall model (the "S" model), or you can choose instead to maintain the saturation current density at the surface at the same value it had in low-level injection:

$$J_o = \frac{q S_n S_p}{S_p \left(p_{eq} + n_{ie} e^{-E_t/kT} \right) + S_n \left(n_{eq} + n_{ie} e^{E_t/kT} \right)}$$

The choice of high-level-injection model for surface recombination has no effect when the surface remains in low-level injection.

Important Note on Surface Recombination

When there is more than one region in the device, the front and back surfaces refer to the current region, not the device as a whole. The recombination rate calculated for the back surface of region 1 is added to the recombination

rate calculated for the front surface of region 2 to get the total rate of recombination at the interface between the two regions.

Device Parameters (Device menu)

This group of selections from the Device menu are parameters that affect the device as a whole, as opposed to influencing only one region.

Area

This command opens a dialog box for specifying the cross-sectional area of the device. The area can be specified in the most convenient units, be that cm^2 , mm^2 , or μm^2 . This area is the projected area of the device, unaffected by surface texture. It is the area over which illumination intensity (watts/cm^2) is figured.

Texture

This selection can be used to open a dialog box for specifying the details of surface texture on the front and/or rear surface of the device. 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 cross-sectional area for current flow is larger at the textured surface than at the back surface by the factor $1/\cos(a)$, where "a" is the facet angle. The facet model assumes that the cross-sectional area decreases with a Gaussian shape from its value at the surface toward the substrate value, with a standard-deviation depth constant equal to one-sixth of the facet depth. 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. Note that (111) facets etched anisotropically into a (100)-oriented surface make a facet angle of 54.74 degrees. The choice of facet angle and depth also affects the photogeneration profile, because it causes photons to travel an oblique path through the device.

Front/Rear Surfaces

This selection opens a dialog box for either the front or rear surface of the device, indicating which of three electrostatic models are to be used. 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 specify the barrier height in electron-volts (eV). A positive value bends the bands downward. 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 ($\text{charges}/\text{cm}^2$). A positive density implies positive external charge, which bends the bands downward, producing a negatively-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.

Circuit Connections (Contacts submenu)

The Circuit connections option opens a dialog box for specifying the physical location of the device's ohmic contacts to the external emitter, base, and collector source circuits. These locations are given as distances from the front surface of the device. If a location is specified that is higher than the device width, then the rear surface of the device is used for that contact. For contacts within the bulk of the device, the precise location of contact may be shifted slightly from that specified so that the contact can be made at a node between finite elements. The default settings are an emitter contact at $x=0$ (the front) and base and collector contact at $x=10,000 \mu\text{m}$ (the rear). You can also specify values for internal series resistance associated with each of the three ohmic contacts. There is a minimum value of these resistances of 1 micro-ohm for numerical-computation reasons. Any voltage drop across these internal series resistances does affect the current-voltage behavior calculated and plotted for the device.

By default, the Emitter and Base contact are Enabled, and the Collector is Disabled. Specifying external source circuits will have no effect unless the corresponding connection to the device for each source has been Enabled. **At least one contact should always be Enabled; otherwise, there is no voltage reference for the device and the solution may not converge.**

The concept of a contact in the middle of a device may seem strange, 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 and the active region of the device, even when the active region is in high-level injection. A small slope in the quasi-Fermi potential between the contact and the active region would result from resistance in the intervening layer. This can be modeled using the internal series resistance. PC1D assumes that all three contacts can be treated as injection points for majority-carrier current at which the majority-carrier quasi-Fermi potential is shifted by the applied voltage, less any drop across the internal resistance.

Internal Shunt Elements (Contacts submenu)

Four internal shunt circuit elements can be enabled (they are normally disabled). Each element can be a conductor, diode, or capacitor. These elements can be connected between any two locations within the device. The precise

location of connection for the anode and cathode may be shifted slightly from the specified positions so that the contact can be made at a node between finite elements. Conductors are specified in terms of their conductance (inverse of resistance). Diodes are specified in terms of a saturation current and an ideality factor. Capacitors are specified in terms of their capacitance.

The internal shunt elements can be used to represent 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 greater. The shunt diode can also be used to model recombination that occurs in an area of the device not explicitly being solved. Attachment of internal shunt elements may adversely impact the convergence speed of the solution, because these elements couple the solution at widely different locations within the device, and this coupling is not fully accounted for in the limited-bandwidth matrix used by PC1D to solve the semiconductor equations.

Reflectance

This selection allows you to specify the external optical reflectance for the front and rear surfaces, and the internal optical reflectances for both surfaces. The internal reflectance can be different for the first encounter of light with that surface than for subsequent encounters. Also, the internal reflectance at each surface can be either specular (polished) or diffuse (lambertian). Internal reflectance at the interfaces between device regions composed of different materials is not accounted for by PC1D.

For the external reflectance of the front or rear surface, a dialog box is opened which gives three options: Fixed, Coated, or External. 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 file containing reflectance data.

To minimize reflectance over a broad spectrum, 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 Si₃N₄ would need to be about 75 nm thick to minimize reflectance at a wavelength of 600 nm. Note that some coating materials may absorb light, especially in the ultra-violet. PC1D does not take this effect into account.

Table 1: Commonly Encountered Indices of Refraction (non-stoichiometric or undensified films will differ)

SiO ₂	silicon dioxide	1.46
Si ₃ N ₄	silicon nitride	2.0
ZnS	zinc sulfide	2.4
MgF ₂	magnesium fluoride	1.4
Al ₂ O ₃	aluminum oxide	1.6
TiO ₂	titanium dioxide	2.5
Ta ₂ O ₅	tantalum pentoxide	2.2
	many organic films	1.4
	window glass	1.5

External reflection files are standard ASCII files with a filename suffix REF. Each line in the file should contain two numerical values, separated by one or more spaces or a tab. The first value on each line is a value of wavelength, in nm. The second value on each line is the reflectance, normalized to unity. The maximum number of lines allowed is 200. 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.

Graph menu commands

The Graph menu provides commands that control which graph is available for inspection using the Interactive-Graph view. It also provides a command for resetting the graphs in the Four-Graph view to their default values, and a command for copying data to the Windows clipboard.

Spatial Graphs	This selection provides a menu of choices for interactive graphs that display spatial information (functions of position). If the information you seek is not included in this list, many more functions are available for a Defined Graph.
Temporal Graphs	This selection provides a menu of choices for interactive graphs that display temporal information (functions of time). Some functions not included in this list can be accessed for a Defined Graph.
Defined Graph	This selection opens a dialog box which allows you to specify x and y functions for any graph, selected from 75 available functions. These functions provide a much more detailed look at the operation of the device than is available from the predefined graphs.
Auxiliary Graph	This command opens a dialog box for selecting a spatial function and a position within the device. Once this information has been entered, any subsequent transient solutions will store the value of that function at that position as a function of time. The data can then be viewed as using the Auxiliary selection from the Temporal Graph menu.
Experimental	This command opens a dialog box for selecting data from an external file. The data can be viewed on a user-defined graph with "Experimental Data" as one of the data types. The experimental data file should be a tab-delimited text file. No unit conversion is performed on this data – it's your responsibility to ensure that the data is sensible.
Default Graphs	This command is only active in Four-Graph view, in which case it resets the four graphs to the default selections. When graphs are chosen for interactive examination, they are added to the four-graph view, displacing the default graphs that were there initially. This is done to make it easy to return to those graphs, by double-clicking on them in the Four-Graph view. However, if your attention shifts to a new issue, you will likely want to restore the default graphs, which are chosen to provide a balanced overview of the device's operation.
Previous History Graph, Next History Graph	Whenever a simulation is completed, the current interactive graph is saved (to a maximum of 100 graphs). You can view these 'history' graphs using the Graph menu or the Page Up/Page Down keys. By flipping through these graphs, you can visually compare simulation results. This feature is especially useful for batch runs.
Retain zoom for history graphs	If this is OFF, the previous (history) graphs will be auto-scaled, so that they are fully visible. Unfortunately, this means that the axes can change, making it hard to compare different results. If Retain Zoom is ON, the axes will not change. You can zoom into a part of the curve, then see how that portion differs from other simulation results. You can zoom out to see the full graph.
Reset history graphs	Removes all the accumulated history graphs, allowing you to start afresh.
Copy Data	This command for copying data to the Windows clipboard for export to other programs is enabled in two situations. First, it is enabled for the Interactive-Graph view, in which case this command copies the data that is currently displayed on the graph. Second, it is enabled in Parameter view if Batch Mode is being used, in which case this command copies the data from the batch table to the clipboard.

Spatial Graphs (Graph menu)

This group of graphs displays information as a function of position in the device, measured as distance from the front surface. Prior to obtaining an equilibrium solution, these graphs display data at each point corresponding to charge-neutral local-equilibrium conditions, with no space-charge region. After an equilibrium solution has been obtained, all of the graphs display physically meaningful values.

Doping Densities	This graph plots the log of donor and acceptor doping densities as a function of position. The solid line is for donors, the dashed line for acceptors. PC1D imposes a minimum non-zero doping density of between 1 and 3 cm ⁻³ for both donors and acceptors. PC1D assumes that all donors and acceptors are ionized.
Carrier Densities	This graph plots the log of electron and hole densities as a function of position. The solid line is for electrons, the dashed line for holes.
Electrostatic Potential	This graph plots the electrostatic potential as a function of position. The zero of potential corresponds to intrinsic Region-1 material in equilibrium. In equilibrium, n-type material has a positive potential, while p-type material has a negative potential.
Electric Field	This graph plots the electric field as a function of position. The electric field in PC1D is uniform within each element and is not usually continuous at the nodes. For this graph, the value shown at each node is a linear interpolation between the values in the two elements adjacent to each node, assuming those values to apply at the mid-point of each element.
Charge Density	This graph plots the charge density as a function of position. The charge density at each node is equal to the hole and donor densities minus the electron and acceptor densities, multiplied by the elementary charge constant.
Current Density	This graph plots the current density (amps per cm ² of cross-sectional area) for electrons and holes. The solid line is for electrons, the dashed line for holes. Also shown, as a dotted line, is the total current density, which is the sum of the electron and hole components. The current density in PC1D may be discontinuous between adjacent elements if there is an interface recombination velocity at a node, or if a circuit element (either external or internal) is connected at the node. In such cases, this graph will display the current densities just to the left of the node; except at $x=0$, where the values shown are the current densities just to the right of that node. Note that at textured surfaces the total current density will not be uniform, because the cross-sectional area for current flow is itself a function of position.
Generation & Recombination	This graph plots the cumulative photogeneration and recombination as a function of position. The solid line is generation, the dashed line is recombination. The values shown are the total integrated quantities from the front surface to each node in the device, including any recombination at that node. Thus, a non-zero value of Recombination at $x=0$ means that there is surface recombination at the front surface. This graph is particularly useful for understanding where the recombination in the device is occurring, since the fraction of the total can be immediately observed by comparing the increase in the plot over a given region to the increase across the entire device. The units in this graph (inverse seconds) can be multiplied by the elementary charge constant to express these quantities in units of current.
Carrier Mobilities	This graph plots the mobilities of electrons and holes as a function of position. The solid line is for electrons, the dashed line for holes. When the variable-mobility model is used, the graph shows the resulting mobility for each carrier at each point based on the doping density at that point and the temperature of the device. Carrier diffusivities can be obtained from this plot by multiplying by the thermal voltage.
Energy Bands	This graph plots the energy-band edges and quasi-Fermi energies as a function of position. The solid line is the conduction-band edge, the

dashed line is the valence-band edge, the dotted line is the electron quasi-Fermi energy, and the dot-dashed line is the hole quasi-Fermi energy. The zero of energy is taken as the equilibrium position of the Fermi energy, and positive energy refers to increasing electron energy (decreasing potential). The quasi-Fermi potentials (and hence energies) are defined at each node and thus are always continuous between adjacent elements.

- Carrier Velocities This graph plots the average group velocity of electrons and holes as a function of position. The solid line is for electrons, the dashed line for holes. Positive velocity is toward the right, away from the front of the device. Current densities are proportional to the product of the velocity times the carrier density. The same conditions apply as described above for Current Density when there is a discontinuity in current at a node.
- Diffusion Length This graph plots the minority-carrier diffusion length as a function of position. This is equal to the square root of the minority-carrier diffusivity times the minority-carrier lifetime, where the lifetime is equal to the excess minority-carrier density divided by the local bulk recombination rate. For equilibrium conditions this ratio would be undefined, so in that case the recombination rate is calculated for a minority carrier density increased by 1% over the equilibrium value.

Temporal Graphs (Graph menu)

This group of graphs displays information as a function of time. While they are primarily useful for transient solutions, they can also be helpful for determining certain values from steady-state calculations, in which case they will display data only for a single point in time (just prior to $t=0$). One of these graphs, Quantum Efficiency, is different from the others in that it displays its data versus illumination wavelength.

Auxiliary	If an auxiliary spatial function has been previously set up using the Auxiliary dialog box accessible from the Graph menu, then this graph will plot that auxiliary data versus time.
Base Current	This graph plots the current into the base connection versus time. The base-contact selections are enabled only if the base contact is enabled and defined for this device.
Base Voltage	This graph plots the voltage at the base connection versus time. The voltage at the node where connection is made is the majority-carrier quasi-Fermi potential at that node. Note that this graph is not equal to the base source voltage when a source series resistance is specified. However, any voltage drop across the internal series resistance of the base connection is incorporated into this graph.
Base I-V	This graph plots the base current versus the base voltage, with each (x,y) pair corresponding to a time step. The left-right cursor keys in parametric graphs such as this correspond to advancing the time steps, not the left-right orientation of the data.
Base I-V / Power	This graph combined the Base I-V data with a plot of base power into the device as a function of base voltage. The ordinate units are mixed, and correspond to amps for current and watts for power. Note that the Base Power is the product of the base current and the base voltage.
Collector Current	Like Base Current, but for the Collector contact. The collector-contact selections are enabled only if the collector contact is enabled and defined for this device.
Collector Voltage	Like Base Voltage, but for the Collector contact.
Collector I-V	Like Base I-V, but for the Collector contact.
Collector I-V / Power	Like Base I-V/Power, but for the Collector contact.
Shunt Element Current	This graph plots the current through each of the four shunt elements versus time. Positive current flow corresponds to current from the anode to the cathode. These elements are defined in the Device Menu using the Contacts:Internal dialog box. This selection is disabled if none of the shunt elements are enabled.
Shunt Element Voltage	This graph plots the voltage across each of the four shunt elements, for the anode with respect to the cathode, versus time. This selection is disabled if none of the shunt elements are enabled.
Quantum Efficiency	This graph plots the spectral performance of the device versus illumination wavelength. This graph functions properly only when the only illumination source is the Primary source and that source is set for a monochromatic spectrum (usually a monochromatic scan). A suitable excitation file for preparing data for this graph is provided with PC1D as scan-qe.exc. There are three curves plotted. The solid line is the internal quantum efficiency, defined as the number of electrons collected per photon not reflected. Reflected photons include both those reflected directly at the incident surface, and those that enter the device but subsequently escape out of the incident surface. The dashed line is the external quantum efficiency, defined as the number of electrons collected per photon incident. The dotted line is the total reflectance, which is the sum of the incident-surface reflectance and the incident-surface escape reflectance.

Defined Graphs (Graph menu)

This dialog box can also be accessed by double-clicking outside the plot area of an interactive graph.

This selection opens a dialog box which allows you to specify x and y functions for the currently selected graph.

Up to four y functions can be plotted against any x function, selected from 75 available functions. The abscissa and ordinate axes can each be scaled either linearly or logarithmically, and you can specify a title for the graph.

Following are descriptions of the available plot functions, separated into Spatial and Temporal quantities. You can plot any spatial function against any other spatial function, or any temporal function against any other temporal function. Experimental data can be plotted against any function.

When you first access this dialog box, it will show the settings for the current Interactive Graph. This makes it easy to add an additional curve to an existing graph. If you'd prefer to start a new graph from scratch, use the radio buttons to toggle between 'Spatial' and 'Temporal' graphs. This will set all curves to '(none)', and will clear the title.

Spatial Functions

Acceptor Doping Density	The net acceptor doping density due to the sum of background and diffused p-type dopants. All of these dopants are considered to be electrically active.
Bulk Recombination Rate	The volumetric recombination rate at each node, including SRH recombination (as influenced by temperature and doping), band-to-band recombination, and Auger recombination.
Charge Density	The volumetric charge density, given by the densities of electrons, holes, donors, and acceptors.
Conduction Band Edge	The effective edge of the conduction band, modified from the actual edge by electrical bandgap narrowing caused by heavy donor doping.
Conductivity	The volumetric conductivity at each node is given by the sum of the product of electron density and mobility plus the product of hole density and mobility. High-field mobility reduction is not included.
Convergence Error	This is the change in normalized potential determined for each node as the result of the previous iteration of the matrix equation solution. The actual change in potential is limited by the potential clamp, which is normally set between 0.1 and 10. The largest value of this error for any node determines the length of the convergence bar located in the status-bar area of the window. The problem is considered solved when the largest error drops below the convergence limit specified in the Compute Numerical dialog box.
Cross-Sectional Area	This is the cross-sectional area at each node through which the electron and hole currents must flow. This will be nonuniform in the vicinity of surface texture.
Cumulative Conductivity	The integrated conductivity, starting at the front surface. This function is zero at the front surface, and at the back surface is equal to the net total sheet conductivity of the device. High-field mobility reduction is not included.
Cumulative Excess Conductivity	The integrated conductivity, starting at the front surface, relative to its value at equilibrium. High-field mobility reduction is not included. At the back surface, this function gives the net total increase in conductivity for the device. This function is useful for simulating photoconductance measurements.
Cumulative Photogeneration	The rate of generation of electron-hole pairs due to external photo-excitation, measured as a cumulative function starting at the front surface of the device. This function is zero at the front surface, and its value at the back surface is the total photogeneration rate for the entire device.

Cumulative Recombination	The recombination rate of excess electrons and holes, calculated as a cumulative function starting at the front surface of the device. In addition to bulk recombination, this function includes surface recombination and interface recombination between regions. Its value at the front surface is the front-surface recombination rate, and its value at the back surface is the total rate of recombination for the entire device.
Dielectric Constant	This is the permittivity of the material at each node, normalized to the permittivity of free space. It is the proportionality constant that relates charge density to electric field.
Diffusion Length	This is the minority-carrier diffusion length, defined as the square root of the diffusivity times minority-carrier lifetime. The lifetime is calculated as the excess carrier density divided by the recombination rate, and the diffusivity is the mobility times the thermal voltage. High-field mobility reduction is considered. In equilibrium, the minority-carrier density is increased by 1% prior to calculating the recombination rate, which includes recombination due to SRH traps, band-to-band, and Auger processes.
Distance from Front	The position of each node, measured as a distance from the front surface. This is the usual function used for the abscissa.
Donor Doping Density	The net donor doping density due to the sum of background and diffused n-type dopants. All of these dopants are considered to be electrically active.
Electric Field	A weighted average of the electric field in the elements to either side of each node (except the end nodes, which are one-sided). The weighting favors the smaller adjoining element, as if the field in each element was assigned to its midpoint, and the resulting value at the interface node calculated as a linear interpolation.
Electron Current	This is the current of electrons past each node.
Electron Current Density	The areal density of electron current flow. Note that since electrons have negative charge, they move in a direction opposite to the sign of their current flow.
Electron Density	The density of mobile electrons.
Electron Diff. Current Density	The electron current density can be broken into drift and diffusion components. The diffusion component is given by the product of the mobility and the gradient of the density of mobile electrons. This function is actually calculated by subtracting the drift component from the total electron current density.
Electron Drift Current Density	The electron current density can be broken into drift and diffusion components. The drift component is given by the product of the mobility, the density of mobile electrons, and the electric field. Each of these terms is determined as they are for plots of these functions, including high-field mobility reduction.
Electron Mobility	The mobility of electrons, which is affected by temperature and doping density of both Donors and Acceptors. It is also reduced by high electric fields which would induce velocity saturation.
Electron Quasi-Fermi Energy	The quasi-Fermi energy for electrons is a non-equilibrium version of the electro-chemical potential which characterizes the statistical thermodynamics of the mobile electrons.
Electron Velocity	The velocity of mobile electrons. This is given by the electron current density divided by the density of mobile electrons.
Electrostatic Potential	The electrostatic potential, where zero equals the potential of intrinsic material in equilibrium. In equilibrium, positive potentials are associated with n-type material, negative with p-type material.
Energy Gap (electrical)	The effective energy gap, which is modified from the material's intrinsic energy gap by electrical bandgap narrowing effects caused by heavy doping. The amount of bandgap narrowing is directly related to the increase in the pn product in equilibrium.

Excess Charge Density	The volumetric charge density at each node measured relative to its value at equilibrium.
Excess Conductivity	The volumetric conductivity at each node relative to its value at equilibrium, including contributions from both electrons and holes. High-field mobility reduction is not included.
Excess Electron Density	The density of mobile electrons in excess of the value at equilibrium.
Excess Electron Density Ratio	This is the density of mobile electrons in excess of the equilibrium value at each node, divided by the equilibrium value at that node. The resulting ratio is dimensionless.
Excess Electrostatic Potential	The electrostatic potential at each node measured relative to its value at equilibrium.
Excess Hole Density	The density of mobile holes in excess of the value at equilibrium.
Excess Hole Density Ratio	This is the density of mobile holes in excess of the equilibrium value at each node, divided by the equilibrium value at that node. The resulting ratio is dimensionless.
Excess pn Product Ratio	The excess pn product is the product of the density of mobile holes and mobile electrons, in excess of the equilibrium value of this product at each node. The ratio is formed by dividing this value by the equilibrium value of the pn product at that node. The resulting ratio is dimensionless.
Generation Rate	The volumetric rate of generation of electron-hole pairs due to external photo-excitation.
Hole Current	This is the current of holes past each node.
Hole Current Density	The areal density of hole current flow.
Hole Density	The density of mobile holes.
Hole Diff. Current Density	The hole current density can be broken into drift and diffusion components. The diffusion component is given by the product of the mobility and the gradient of the density of mobile holes. This function is actually calculated by subtracting the drift component from the total hole current density.
Hole Drift Current Density	The hole current density can be broken into drift and diffusion components. The drift component is given by the product of the mobility, the density of mobile holes, and the electric field. Each of these terms is determined as they are for plots of these functions, including high-field mobility reduction.
Hole Mobility	The mobility of holes, which is affected by temperature and doping density of both Donors and Acceptors. It is also reduced by high electric fields which would induce velocity saturation.
Hole Quasi-Fermi Energy	The quasi-Fermi energy for holes is a non-equilibrium version of the electro-chemical potential which characterizes the statistical thermodynamics of the mobile holes.
Hole Velocity	The velocity of mobile holes. This is given by the hole current density divided by the density of mobile holes.
Intrinsic Conc. (effective)	The square-root of the pn product in equilibrium. Nominally, this equals the concentration of electrons and holes in intrinsic material held in equilibrium, but heavy doping can increase the pn product, and hence the effective intrinsic carrier concentration.
Minority Carrier Lifetime	This is the excess minority carrier density divided by the recombination rate. In equilibrium this ratio is undefined, in which case the minority carrier concentration is increased by 1% before calculating the recombination rate. The minority carrier is whichever carrier (electrons vs. holes) currently has a lower density. If the lifetime calculated this way is ever negative, it is replaced with zero.
Normalized Excess	The excess pn product is the product of the density of mobile holes and

pn Product	mobile electrons, in excess of the equilibrium value of this product at each node. A ratio is formed by dividing this value by the equilibrium value of the pn product at that node. The function is then normalized by dividing the value at each node by the largest value found in the device. The resulting ratio is dimensionless and has a maximum value of unity. This function when plotted for forward bias without photoexcitation is equal to the probability that a carrier pair photogenerated at that location would be collected by the external circuit under short-circuit illumination conditions.
Resistivity	The resistivity of the material at each node is equal to the inverse of the conductivity, which is given by the sum of the product of electron density and mobility plus the product of hole density and mobility. High-field mobility reduction is not included in this function.
Total Current	This is the sum of currents due to electrons and holes. In steady state, this function should be uniform between external contacts. During transients, nonuniformities will emerge in regions where a net charge density is accumulating or depleting.
Total Current Density	The sum of the electron and hole current densities.
Vacuum Energy	The energy an electron would have if a void were formed in the device at that point and the electron were to escape into that void. It is equal to the energy of the conduction band edge plus the electron affinity of the material.
Valence Band Edge	The effective edge of the valence band, modified from the actual edge by electrical bandgap narrowing caused by heavy acceptor doping.

Temporal Functions

Auxiliary	If an auxiliary spatial function has been previously set up using the Auxiliary dialog box accessible from the Graph menu, then this function will plot that auxiliary data.
Base Current	The current into the base contact to the device. This current is injected as majority-carrier current at the node nearest to the specified contact location.
Base Power	The product of the voltage at the base contact times the current into this contact. This power is positive for energy flow into the device, and negative for energy flow out of the device.
Base Voltage	The voltage at the base contact to the device. This differs from the base source voltage when a series source resistance has been specified in the Excitation BaseCircuit dialog. It differs from the majority-carrier quasi-Fermi potential at the location of the base contact only by an amount equal to the voltage drop across the internal series resistance associated with the contact, as specified in the Device Circuit Contacts dialog.
Collector Current	The current into the collector contact to the device. This current is injected as majority-carrier current at the node nearest to the specified contact location.
Collector Power	The product of the voltage at the collector contact times the current into this contact. This power is positive for energy flow into the device, and negative for energy flow out of the device.
Collector Voltage	The voltage at the collector contact to the device. This differs from the collector source voltage when a series source resistance has been specified in the Excitation CollectorCircuit dialog. It differs from the majority-carrier quasi-Fermi potential at the location of the collector contact only by an amount equal to the voltage drop across the internal series resistance associated with the contact, as specified in the Device Circuit Contacts dialog.
Elapsed Time	The total elapsed time since the start of the transient solution. This is

the usual value for the abscissa in temporal graphs. The steady-state solution is represented as occurring at negative time, with the start of the transient at $t=0$.

External Quantum Efficiency	This function plots the rate that charge carriers are collected by the base circuit divided by the rate that photons are incident on the device from the primary source. To properly represent EQE, the problem must have been set up with the base circuit at short circuit and the primary illumination set to scan a range of monochromatic wavelengths.
Internal Quantum Efficiency	This function is similar to the External Quantum Efficiency, except that the result is divided by $1-R$, where R is the sum of the Pri-Surface Reflectance and Pri-Surface Escape. IQE represents the rate at which carriers are collected by the external circuit divided by the net rate that photons enter the surface of the device. Here, "net rate" means the rate of incidence minus the rate of exit. If a secondary light source is present, use 'IQE adjusted for light bias' instead.
Inverse IQE	This function simply plots the inverse of the Internal Quantum Efficiency. Plots of inverse IQE versus Absorption Length are commonly used to extract information about the diffusion length in the device.
IQE adjusted for light bias	This is the Internal Quantum Efficiency, but remains valid when the device is being illuminated by a secondary (constant) light source.
Pri-Surface Absorption Length	This function plots the absorption length in the material for monochromatic illumination from the primary source. When the device is composed of multiple materials, the material at the surface onto which the primary source is incident is used. This function is defined only when the problem is set up with primary source illumination configured to scan a range of monochromatic wavelengths.
Pri-Surface Escape	This function plots the fraction of incident photons that enter the device initially, but which are then reflected internally and escape from the surface of the device onto which the primary source illumination is incident. It is defined only when the problem is set up with primary source illumination configured to scan a range of monochromatic wavelengths.
Pri-Surface Reflectance	This function plots the fraction of incident photons that are reflected directly from the surface of the device onto which the primary source illumination is incident. It is defined only when the problem is set up with primary source illumination configured to scan a range of monochromatic wavelengths.
Pri-Surface Refractive Index	This function plots the refractive index of the material for monochromatic illumination from the primary source. When the device is composed of multiple materials, the material at the surface onto which the primary source is incident is used. This function is defined only when the problem is set up with primary source illumination configured to scan a range of monochromatic wavelengths.
Pri-Surface Total Reflectance	This is just Pri-Surface Reflectance + Pri-Surface Escape.
Primary Source Wavelength	The wavelength (in free space) of light incident on the device from the primary source. This function is defined only when monochromatic illumination has been specified.
Shunt ## Current	This is a set of four function that plot the current through the four internal elements defined using Device Circuit InternalElements. The sign of the current corresponds to current entering the element at its anode and leaving at its cathode.
Shunt ## Voltage	This is a set of four functions that plot the voltage across the four internal elements defined using Device Circuit InternalElements. The sign of the voltage corresponds to the voltage of the element anode with respect to its cathode.

Bibliography

PC1D Program Development:

Clugston, D.A. and Basore, P.A., "PC1D Version 5: 32-bit Solar Cell Simulation on Personal Computers," 26th IEEE Photovoltaic Specialists Conf. (Sept 1997).

Basore, P.A. and Clugston, D.A., "PC1D Version 4 for Windows: From Analysis to Design," 25th IEEE Photovoltaic Specialists Conf., (May 1996).

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).

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)

Fan, Y.H., "Semiconductors and Semimetals", ed. R.K. Williamson and A.C. Beer, Academic Press (1967), Vol 3, p.409. (III-V free-carrier absorption)

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)

Hurkx, Klaassen, and Knuvers, "A New Recombination Model for Device Simulation Including Tunneling", IEEE Trans. Electron Devices, ED-39, (Feb. 1992), pp. 331-338. (Field-enhanced recombination).

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)
- Schlengenotto, Maeder, and Gerlach, "Temperature Dependence of the Radiative Recombination Coefficient in Silicon", Physica Status Solidi A21, 357-367 (1974). (Silicon recombination)
- Schroder, Thomas, and Swartz, "Free Carrier Absorption in Silicon", IEEE Trans. Electron Devices, ED-25, (Feb. 1978), pp. 254-261. (Si free-carrier absorption)
- 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. (AM0 Spectrum)
- Thurber, Mattis, Liu, and Filliben, "The Relationship Between Resistivity and Dopant Density for Phosphorus- and Boron-Doped Silicon," U.S. Department of Commerce National Bureau of Standards, (1981). (Si majority carrier mobilities)
- 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, Misiakos, and Neugroschel, IEEE Trans. on Electron Devices ED-37 (5), May 1990, p. 1314. (Silicon minority hole mobility).
- Wiley, J., Semiconductors and Semimetals, Vol. 10, Academic Press, N.Y. (1975). (InP, GaAs, AlGaAs transport 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.

NUMERICAL METHOD

This section describes the isolated-element approach used within PC1D to obtain rapid, accurate solutions of the semiconductor transport equations. 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 A.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 PC1D 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.

A.1 Fundamental Equations

The numerical method of PC1D 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}).

$$\begin{aligned} J_n &= \mu_n n \nabla E_{Fn} \\ J_p &= \mu_p p \nabla E_{Fp} \end{aligned} \quad (\text{A.1})$$

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 non-degenerate material, the Boltzmann exponential approximation to the Fermi distribution function can be utilized:

$$\begin{aligned} n &= N_c e^{-(E_c - E_{Fn})/kT} \\ p &= N_v e^{-(E_{Fp} - E_v)/kT} \end{aligned} \quad (\text{A.2})$$

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 PC1D, this location is taken as $x=0$. The position of the two band edges in Eq. (A.2) 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 centre 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}),

$$\begin{aligned} V_n &= \left(\frac{kT}{q} \right) \ln \left(\frac{N_c}{N_{cr}} \right) + \frac{\chi - \chi_R}{q} + \Delta E_{gc} \\ V_p &= \left(\frac{kT}{q} \right) \ln \left(\frac{N_v}{N_{vr}} \right) - \frac{E_g - E_{gR}}{q} - \frac{\chi - \chi_R}{q} + \Delta E_{gv} \end{aligned} \quad (\text{A.3})$$

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. (A.2) produces a new pair of Boltzmann-like expressions for n and p ,

$$\begin{aligned} n &= n_{ir} e^{q(\psi + V_n - \phi_n) / kT} \\ p &= n_{ir} e^{q(-\psi + V_p + \phi_p) / kT} \end{aligned} \quad (\text{A.4})$$

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

Combining Eqs. (A.1) and (A.4) 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).

$$\begin{aligned} \frac{\partial n}{\partial t} &= \frac{\nabla \cdot \mathbf{J}_n}{q} + G_L - U_n \\ \frac{\partial p}{\partial t} &= \frac{\nabla \cdot \mathbf{J}_p}{q} + G_L - U_p \end{aligned} \quad (\text{A.5})$$

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 PC1D includes a number of sophisticated features, including light trapping and free-carrier absorption. It does not consider 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. PC1D 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. 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. (A.5) 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.

$$\nabla \cdot (\epsilon \nabla \Psi) = -\rho \quad (\text{A.6})$$

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. PC1D 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 PC1D 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. PC1D 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, PC1D 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, PC1D

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 modelled using the external series resistance.

A.2 Discretization

The first step in the numerical solution of the semiconductor equations described in Section A.1 is the discretization of the device, both in time and space. PC1D 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 .

$$W_e = \int_0^w \frac{\epsilon}{2} \left[\frac{\partial \psi}{\partial x} \right]^2 dx - \int_0^w \rho \psi dx \quad (A.7)$$

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:

$$\epsilon^- \frac{A_{k-1} + A_k}{2} \frac{\Delta \psi_{k-1}}{\Delta x_{k-1}} - \epsilon^+ \frac{A_{k-1} + A_k}{2} \frac{\Delta \psi_{k-1}}{\Delta x_{k-1}} = \frac{\rho_{k-1} \Delta x_{k-1} A_{k-1} + 2\rho_k (\Delta x_{k-1} + \Delta x_k) A_{k-1} + \rho_{k+1} \Delta x_k A_{k+1}}{6} \quad (A.8)$$

where ϵ^- and ϵ^+ refer to the permittivity of the elements on the right and left sides of the node, respectively. In Eq. (A.8), 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. (A.1), which gives the electron current, I_n , in the direction of the solution dimension.

$$I_n(x) \cdot \mathbf{a}_x = -q \mu_n(x, E) n(x) A(x) \nabla \phi_n(x) \quad (\text{A.9})$$

The mobility in Eq. (A.9) is electric-field dependent. The model used for this dependency is

$$\mu(x, E) = \mu(x, 0) / \sqrt{1 + [\mu(x, 0)E / v_s(x)]^2} \quad (\text{A.10})$$

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. (A.10) is assumed to be uniform within each element.

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

$$\xi_n(x) = \psi(x) + V_n(x) + \frac{k}{q} \frac{T \mu_n(x, E) A(x)}{\mu_r A_r} \quad (\text{A.11})$$

where μ_r and A_r are arbitrary fixed reference values for mobility and area. Substitute the potentials of Eq. (A.4) for n in Eq. (A.9) and rewrite the expression in terms of the modified potential of Eq. (A.11).

$$I_n(x) \mathbf{a}_x = -q \mu_n n_i \frac{A_r}{A} \nabla \phi_n(x) e^{-\frac{q}{k} [\xi_n(x) - \phi_n(x)]} \quad (\text{A.12})$$

Eq. (A.12) can be rearranged so that the quasi-Fermi potential is the only position-dependent value on the RHS of the equation.

$$I_n(x) e^{-q\xi_n(x)/k} \mathbf{a}_x^T = q n_i \frac{A_r}{A} \nabla e^{-q\phi_n(x)/k} \quad (\text{A.13})$$

where $D_r = \mu_r kT/q$ 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} .

$$\int_c I_n(x) e^{-q\xi_n(x)/k} \mathbf{a}_x^T \cdot \mathbf{d} = q n_i \frac{A_r}{A} \int_c \nabla e^{-q\phi_n(x)/k} \cdot \mathbf{d} \quad (\text{A.14})$$

$$\int_{x_k}^{x_{k+1}} I_n(x) e^{-q\xi_n(x)/k} dx = q n_i \frac{A_r}{A} \left[e^{-q\phi_n(x_k)/k} - e^{-q\phi_n(x_{k+1})/k} \right] \quad (\text{A.15})$$

The value of this approach is now apparent. To complete the integration on the LHS of Eq. (A.15), it is only necessary to make assumptions regarding the spatial variation of two well-behaved variables: I_n and xn . No assumption regarding the spatial variation of n was necessary. Since the carrier concentration, n , depends on both xn 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 xn 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:

$$I_n(x_k) = -\Delta I_n(\beta \Delta \xi_n) + \frac{q_r n_i D_r}{\Delta x} \left[Z(\beta \Delta \xi_n) e^{-\beta(\xi_n(x_{k+1}) - \phi_n(x_{k+1}))} - Z(-\beta \Delta \xi_n) e^{-\beta(\xi_n(x_k) - \phi_n(x_k))} \right] \quad (\text{A.16})$$

where $Z(u) = u / (eu - 1)$ and $Y(u) = (1 - Z(u)) / u$

In Eq. (A.16), 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 PC1D, 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:

$$I_p(x_k) = -\Delta I_p Y(-\beta \Delta \xi_p) - \frac{q_r n_i D_p}{\Delta x} \left[Z(\beta \Delta \xi_p) e^{-\beta(\xi_p(x_{k+1}) - \phi_p(x_{k+1}))} - Z(-\beta \Delta \xi_p) e^{-\beta(\xi_p(x_k) - \phi_p(x_k))} \right] \quad (\text{A.17})$$

A.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 A.2, two equations at each internal node k result.

Hole Current Transport:

$$\begin{aligned}
& -\frac{D_r}{\Delta x_{k-1}} \left[Z(-\beta \Delta \xi_p)_{k-1} e x \mathfrak{P} \mathfrak{B}(\xi_p - \phi_p)_{k-1} - Z(\beta \Delta \xi_p)_{k-1} e x \mathfrak{P} \mathfrak{B}(\xi_p - \phi_p)_{k-1} \right] \\
& + \frac{(\Delta I_p)_{k-1}}{n_{ir} A_r} Y(\beta \Delta \xi_p)_{k-1} \\
& = -\frac{D_r}{\Delta x_k} \left[Z(-\beta \Delta \xi_p)_k e x \mathfrak{P} \mathfrak{B}(\xi_p - \phi_p)_{k+1} - Z(\beta \Delta \xi_p)_k e x \mathfrak{P} \mathfrak{B}(\xi_p - \phi_p)_k \right] \\
& - \frac{(\Delta I_p)_k}{n_{ir} A_r} Y(-\beta \Delta \xi_p)_k + \frac{R_s}{n_{ir}} - \frac{I_{m a j}}{q n_r A_r}
\end{aligned} \tag{A.18}$$

Electron Current Transport:

$$\begin{aligned}
& \frac{D_r}{\Delta x_{k-1}} \left[Z(\beta \Delta \xi_n)_{k-1} e x \mathfrak{P} \mathfrak{B}(\xi_n - \phi_n)_k - Z(-\beta \Delta \xi_n)_{k-1} e x \mathfrak{P} \mathfrak{B}(\xi_n - \phi_n)_{k-1} \right] \\
& + \frac{(\Delta I_n)_{k-1}}{n_{ir} A_r} Y(-\beta \Delta \xi_n)_{k-1} = \\
& \frac{D_r}{\Delta x_k} \left[Z(\beta \Delta \xi_n)_k e x \mathfrak{P} \mathfrak{B}(\xi_n - \phi_n)_{k+1} - Z(-\beta \Delta \xi_n)_k e x \mathfrak{P} \mathfrak{B}(\xi_n - \phi_n)_k \right] \\
& - \frac{(\Delta I_n)_k}{n_{ir} A_r} Y(\beta \Delta \xi_n)_k + \frac{R_s}{n_{ir}} - \frac{I_{m a j}}{q n_r A_r}
\end{aligned} \tag{A.19}$$

Eqs. (A.18) and (A.19), combined with Eq. (A.8), provide three simultaneous non-linear 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. (A.18) and (A.19), with the expression for current that extends beyond the boundaries of the device set to zero (ie. 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 non-linear 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. (A.20).

$$\begin{bmatrix} \frac{\partial f_1}{\partial \phi_n^-} & \frac{\partial f_1}{\partial \psi^-} & \frac{\partial f_1}{\partial \phi_p^-} & \frac{\partial f_1}{\partial \phi} & \frac{\partial f_1}{\partial \psi^0} & \frac{\partial f_1}{\partial \phi_p^0} & \frac{\partial f_1}{\partial \phi_n^+} & \frac{\partial f_1}{\partial \psi^+} & \frac{\partial f_1}{\partial \phi_p^+} \\ \frac{\partial f_2}{\partial \phi_n^-} & \frac{\partial f_2}{\partial \psi^-} & \frac{\partial f_2}{\partial \phi_p^-} & \frac{\partial f_2}{\partial \phi} & \frac{\partial f_2}{\partial \psi^0} & \frac{\partial f_2}{\partial \phi_p^0} & \frac{\partial f_2}{\partial \phi_n^+} & \frac{\partial f_2}{\partial \psi^+} & \frac{\partial f_2}{\partial \phi_p^+} \\ \frac{\partial f_3}{\partial \phi_n^-} & \frac{\partial f_3}{\partial \psi^-} & \frac{\partial f_3}{\partial \phi_p^-} & \frac{\partial f_3}{\partial \phi} & \frac{\partial f_3}{\partial \psi^0} & \frac{\partial f_3}{\partial \phi_p^0} & \frac{\partial f_3}{\partial \phi_n^+} & \frac{\partial f_3}{\partial \psi^+} & \frac{\partial f_3}{\partial \phi_p^+} \end{bmatrix} \begin{bmatrix} \Delta \phi_n^0 \\ \Delta \psi^0 \\ \Delta \phi_p^0 \end{bmatrix} = - \begin{bmatrix} f_1 \\ f_2 \\ f_3 \end{bmatrix} \quad (\text{A.20})$$

where superscripts -, o, 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 $5 kT/q$.

$$\Delta \psi' = \Delta \psi / (1 + |\Delta \psi / \text{clamp}|) \quad (\text{A.21})$$

The changes in the two quasi-Fermi potentials at each node are also limited, in one of three ways selected by the user. In *Phi* clamping, equation (A.21) is applied to each quasi-Fermi potential. In *Psi* clamping, the separation between each quasi-Fermi potential and the electrostatic potential is monitored and prevented from changing by more than the clamp value during each iteration, using equation (A.22). Applying *Both* clamps usually gives the most reliable convergence, through there are exceptions.

$$\Delta \phi' = \Delta \psi' + (\Delta \phi - \Delta \psi') / (1 + |(\Delta \phi - \Delta \psi') / \text{clamp}|) \quad (\text{A.22})$$

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 $1.0\text{E-}6$.

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 non-linear, and the same iterative Newton method is applied. At the end of each iteration ψ is updated, with the clamping function of Eq. (A.21) applied. The iterations continue until the maximum change in $\beta \psi$ is less than the user-supplied normalized error limit.

A.4 Node Placement

The flexibility of PC1D 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. After renoding, the error will usually increase for a few iterations, then converge rapidly toward zero.

The node placement algorithm used in PC1D 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 500 elements are necessary to achieve the specified level of accuracy, PC1D will automatically

increase the element size factor until the number of elements drops below 500. The default value for this factor (α) is 0.5. For more accuracy, consider using a smaller value of α , like 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, PC1D 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:

$$\left(\frac{\rho \Delta x^2}{2\epsilon} \right)^{1/2} > \alpha \frac{k T}{q} \quad (\text{A.23})$$

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

$$| \Delta \ln(N_{\text{maj}}) | > \alpha \quad (\text{A.24})$$

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

$$\Delta x > x. \quad (\text{A.25})$$

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

$$\Delta x > w_{\text{reg}}/20. \quad (\text{A.26})$$

- (1) 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 are 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 , x_n , x_p , v_{sat} , E_t , ϵ