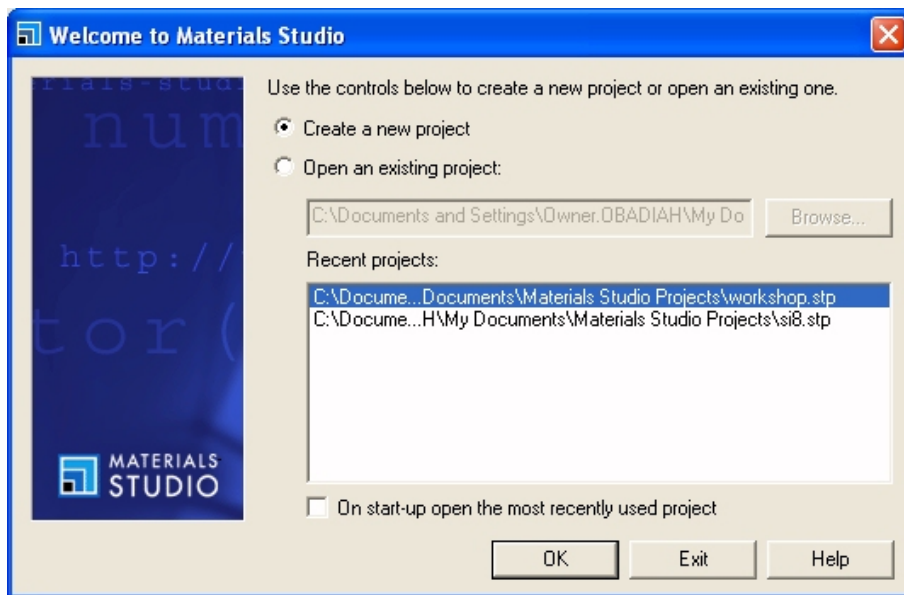


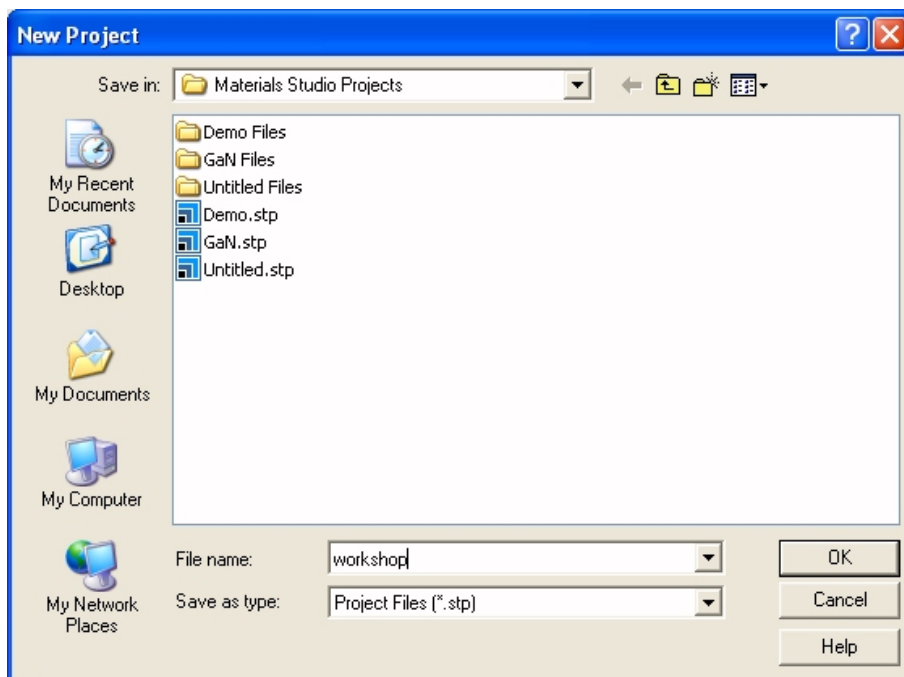
A Beginner's Guide to Materials Studio and DFT Calculations with Castep

P. Hasnip (pjh503@york.ac.uk)

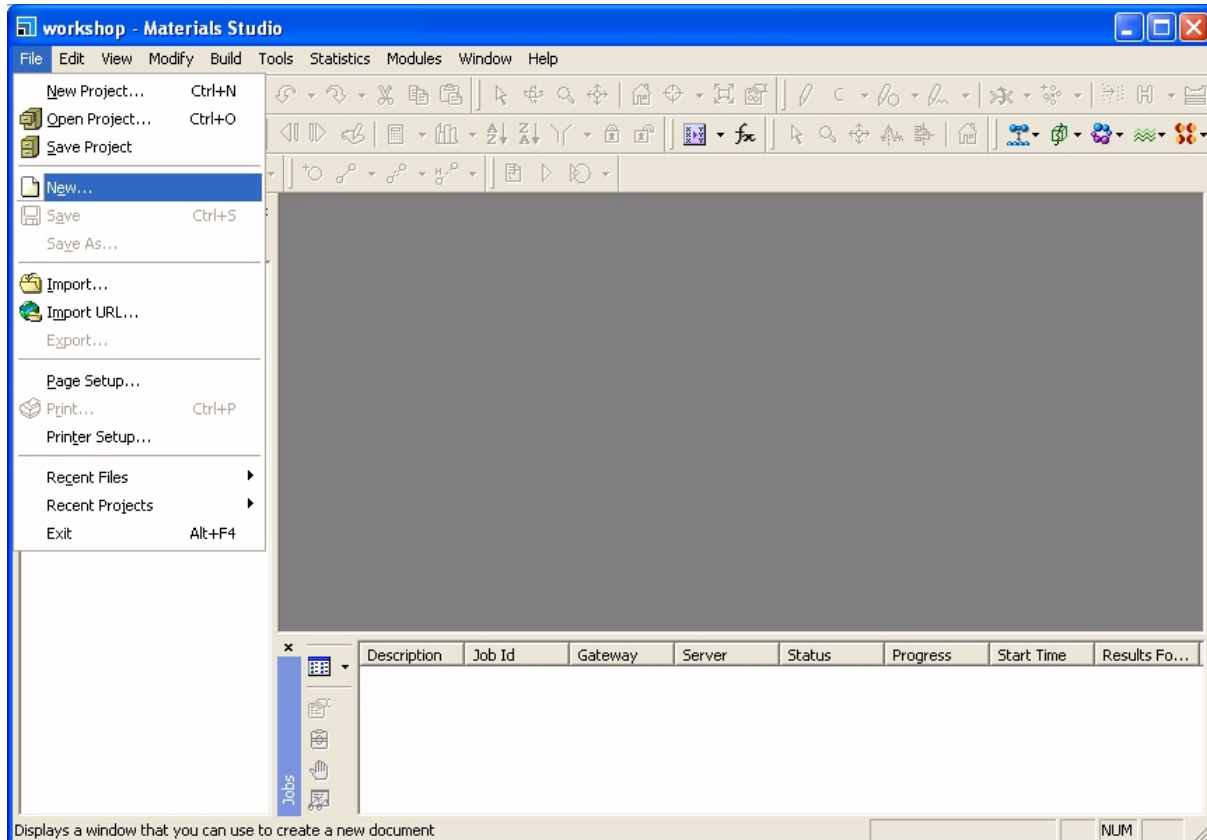
September 18, 2007

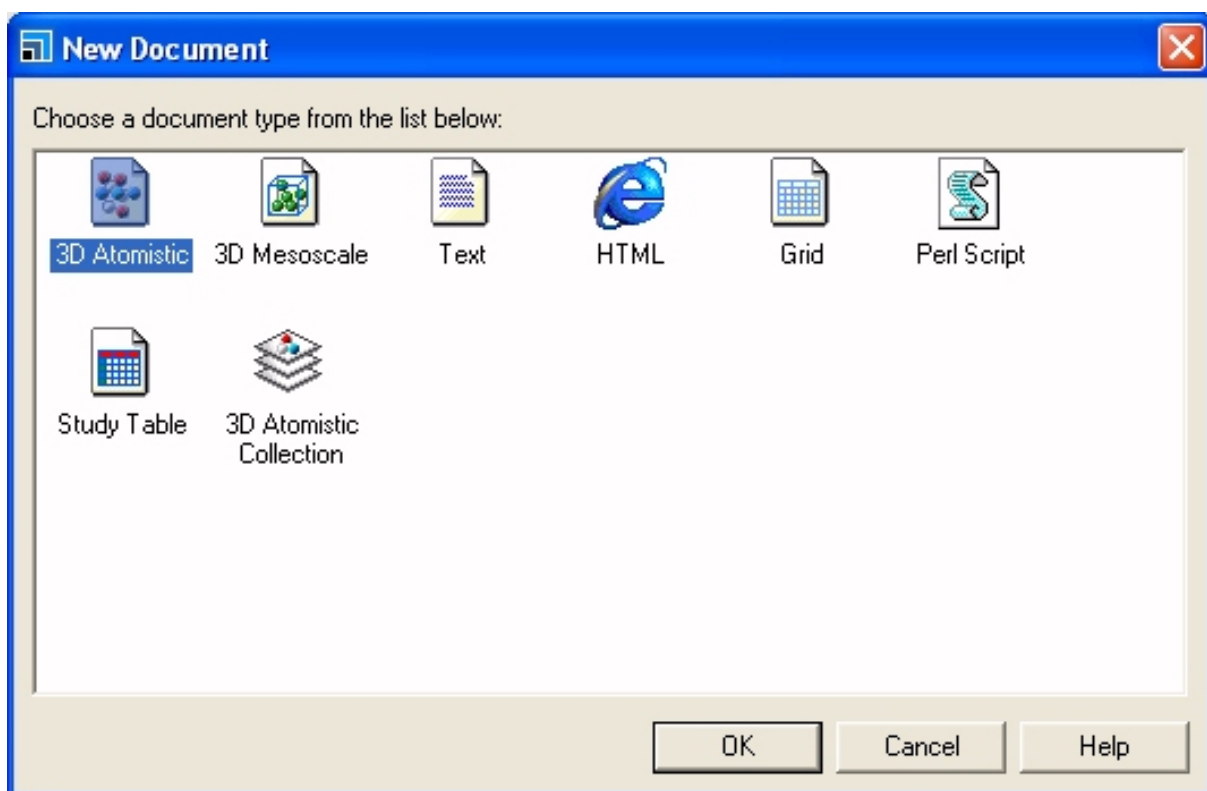


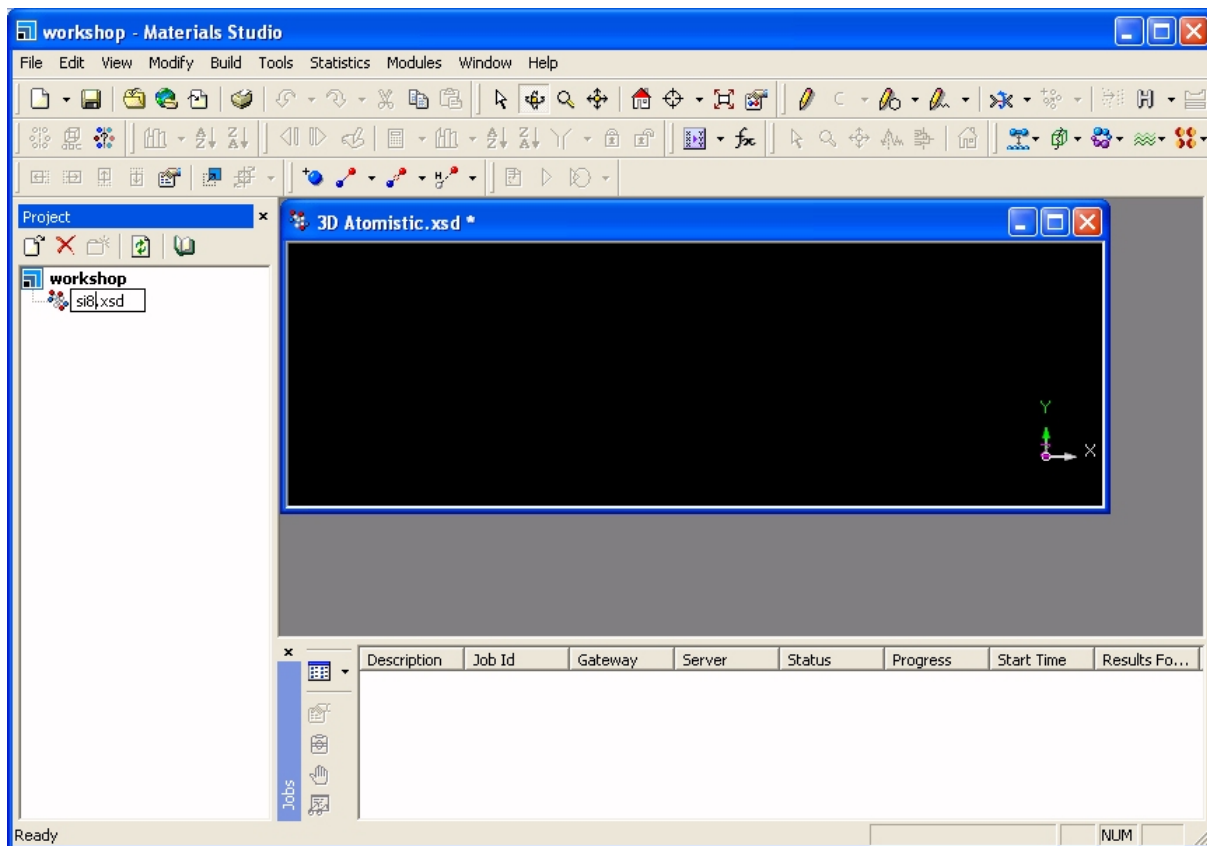
Materials Studio collects all of its files into “Projects”. We’ll start by creating a new project.



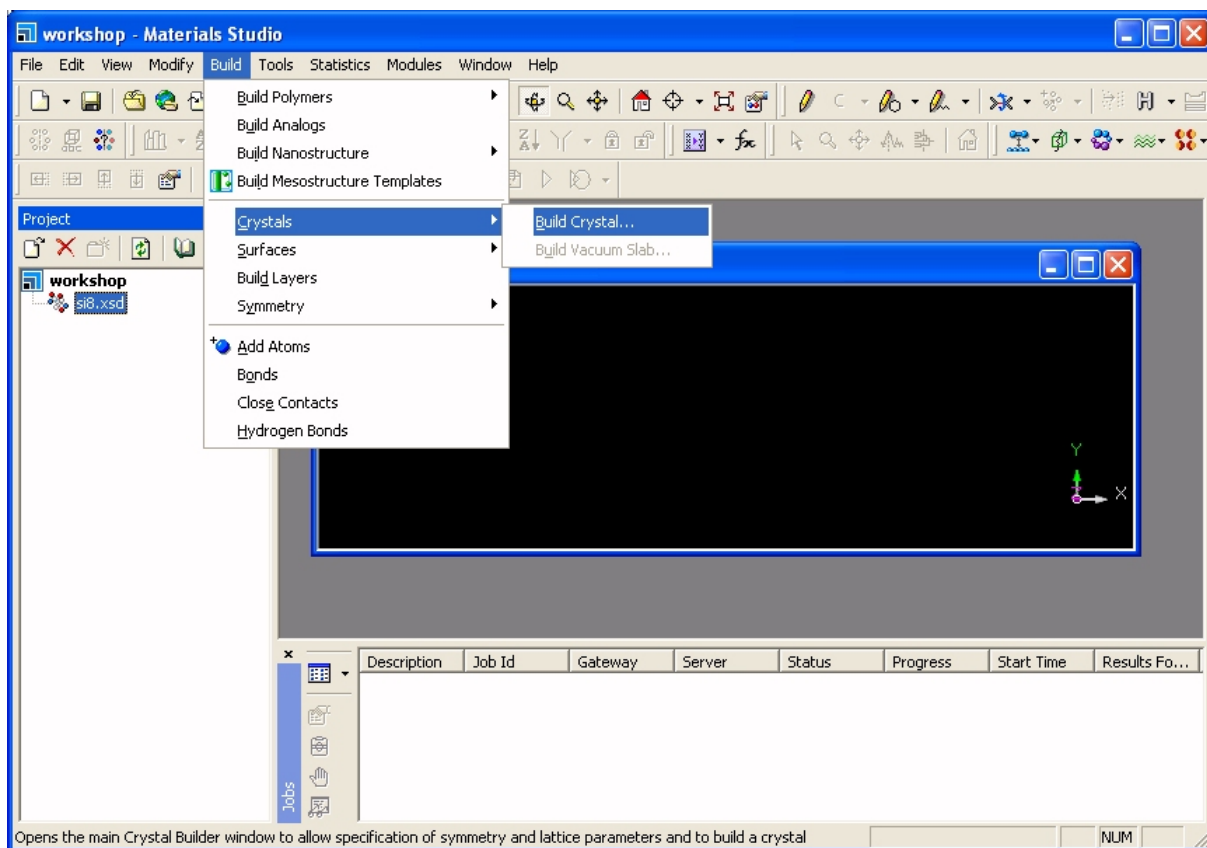
Now we've got a blank project, and we want to define a simulation cell to perform a Castep calculation on. First we add a "3D Atomistic document".

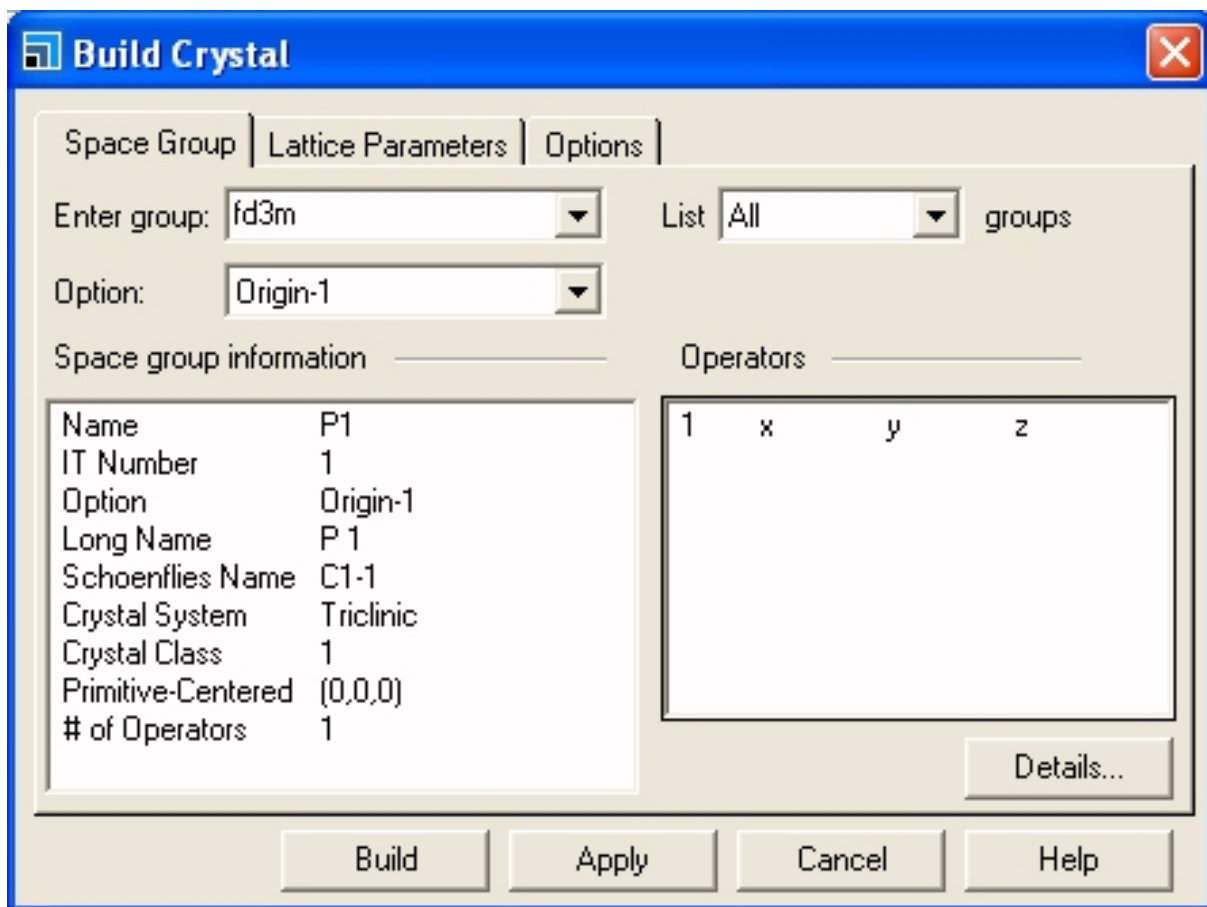






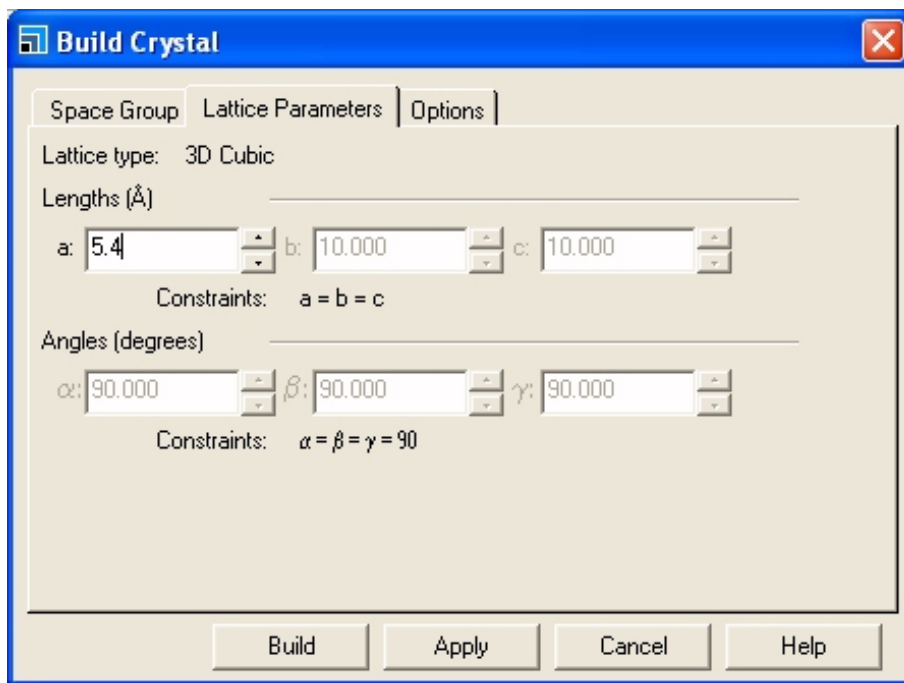
We're going to start by simulating an eight atom silicon FCC cell, so rename the file accordingly. First we'll create the unit cell.





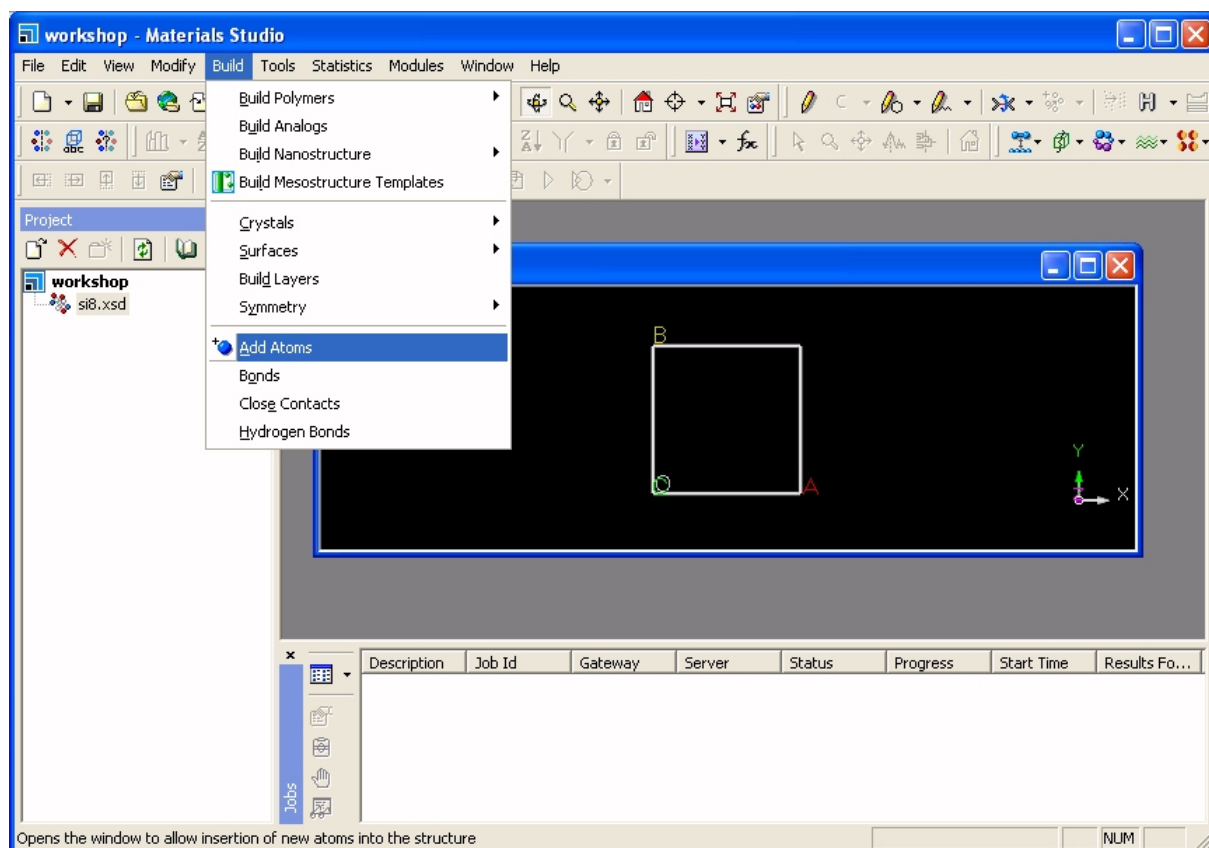
The default is space group P1, i.e. no symmetry. Silicon has the diamond structure (space group FD3M). By telling Materials Studio this symmetry it will automatically apply it to the atoms, thus generating atoms at the symmetry points.

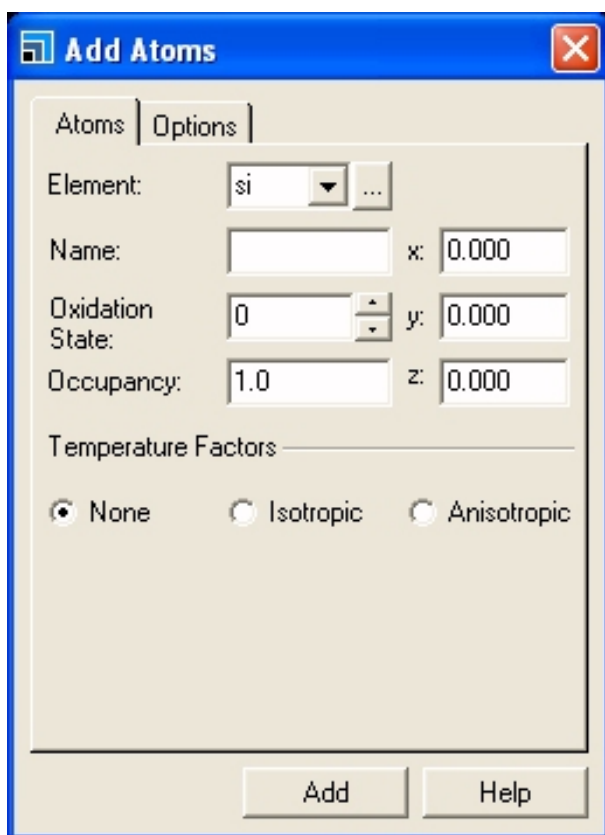
Now to add the lattice constant – click on the “Lattice” tab near the top of the “Build Crystal” window. Since FD3M is cubic (FCC) Materials Studio knows only a has to be set, and the angles and other lattice constants are greyed-out.



Enter “5.4”, and then click on “Build”.

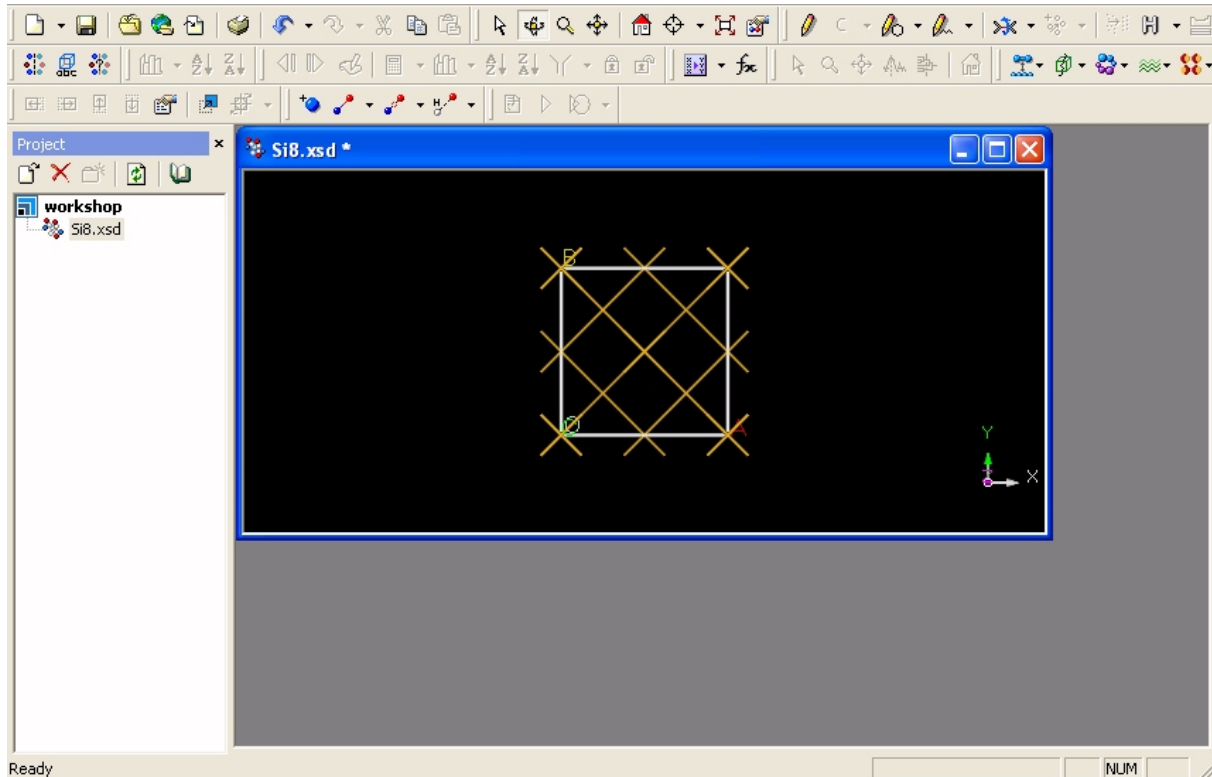
Now we'll add a single silicon atom...





Add a silicon atom at the origin, by changing the “element” from its default and clicking “Add”. By default the co-ordinates are in fractionals, but you can change this on the “Options” tab.

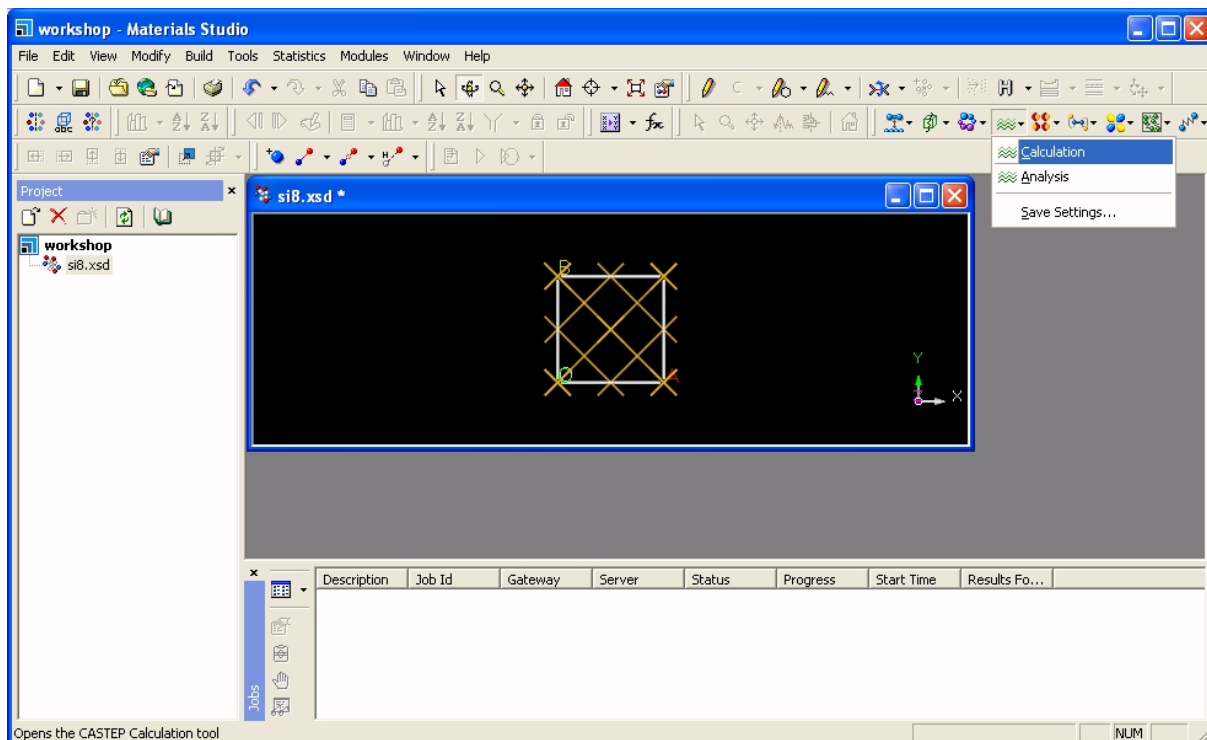
Since we'd already told Materials Studio what the symmetry of the crystal was, our single silicon atom is replicated at each symmetry site and we now have a shiny new eight-atom silicon unit cell.



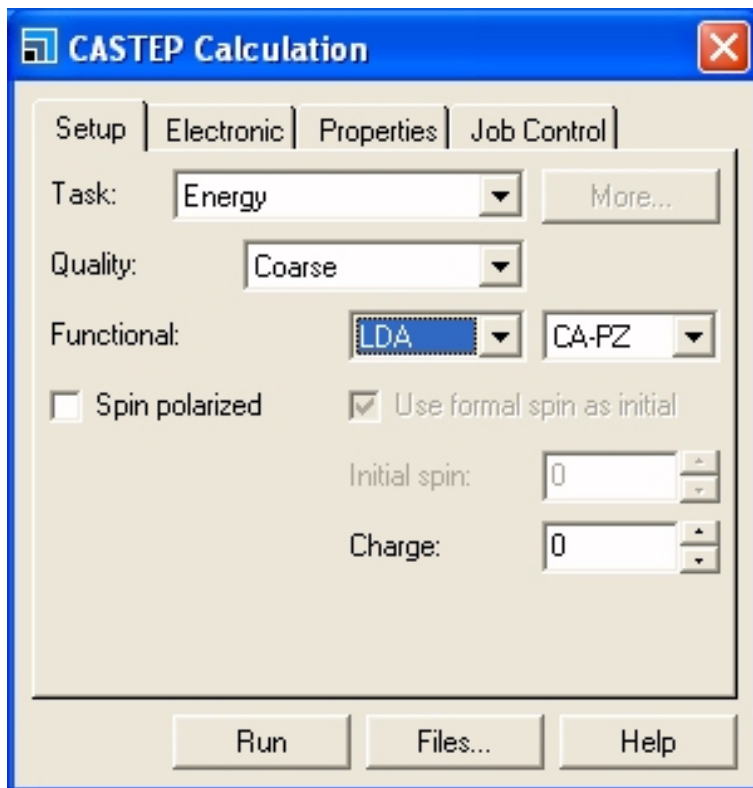
You can rotate the view by holding down the left or right mouse button and dragging, or move it by holding down the middle button. Use the mouse wheel, or both the left and right buttons simultaneously, to zoom in and out.

By default the atoms are shown as little crosses with lines for bonds, and silicon atoms are coloured brownish orange. You can always change this if you don't like it. The “bonds” are just guesses made by Materials Studio based on the element's typical bond-lengths.

We're now ready to run Castep to find the groundstate charge density. Click on the Castep icon, which is a set of three wavy lines (to represent plane-waves), and select “Calculation”.

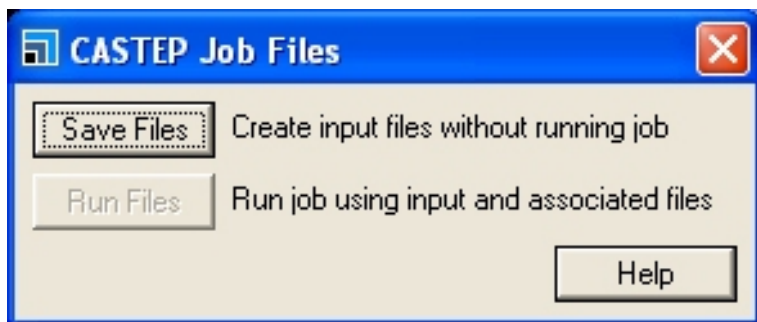


Materials Studio offers a high-level interface to Castep, with cut-off energy, k-point sampling, convergence tolerances etc. all set by the single setting “Quality”. We’ll look at how to specify these things later, but for now we’ll just do a very quick, rough calculation of the groundstate energy and density of our cell.




Make sure the task is “Energy”, and select “Coarse” for quality, and “LDA” for the XC functional.

If you want to run your calculations on Lagavulin (or anywhere else you have Castep available) then you'll want to click "Files" in the Castep window.



Simply select "Save Files" to save the cell and param files. By default these are written to a folder in "My Documents" called "Materials Studio Projects" but be warned – cell files are hidden files, and you won't be able to see them unless you tell Windows you want to view "Hidden and System Files" for that folder.

If you want to run Castep on the PC you're using, you just need to click "Run" on the Castep window. You should see this window appear:



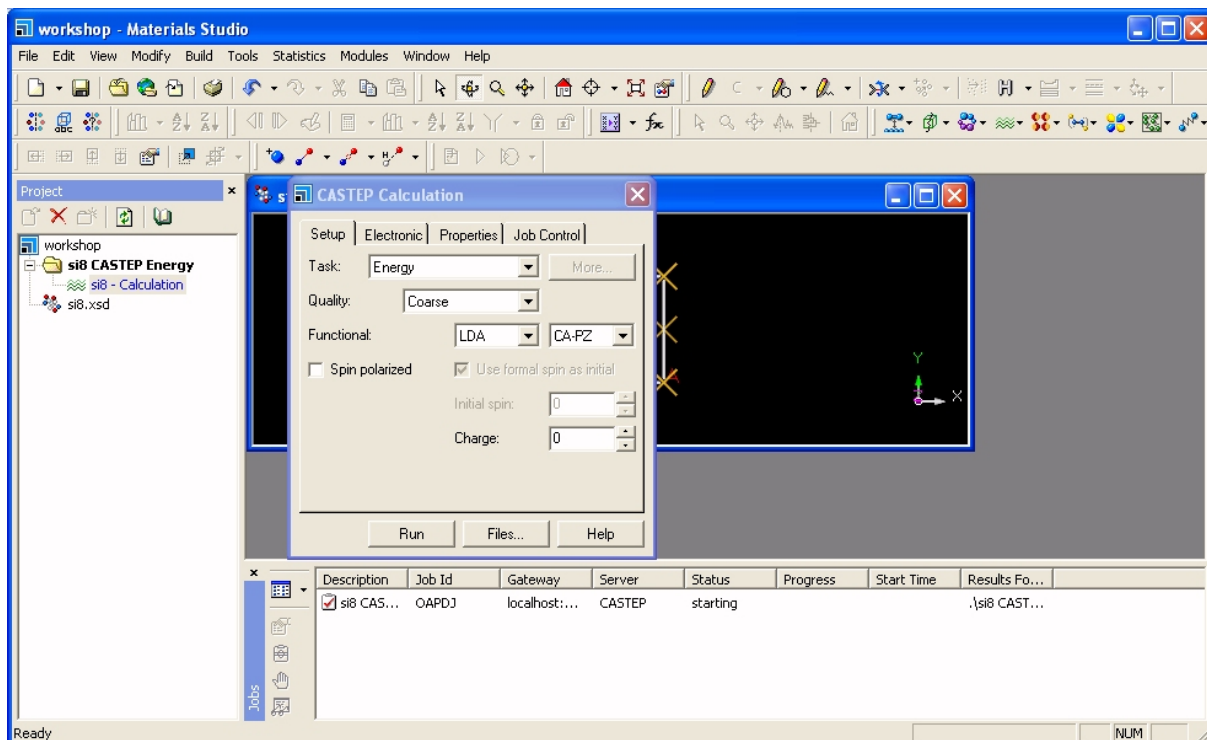
Materials Studio is telling you that your system isn't actually the primitive unit cell, and it's offering to convert it to the primitive cell for you. For now choose "No".

Castep runs via a "Gateway", which might be on your local computer or on a remote machine. This Gateway handles Materials Studio's requests to run calculations and copies the files to and from the "Castep Server".

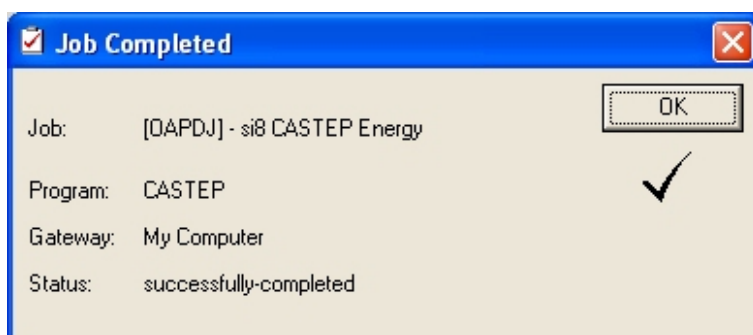
Since the Gateway is actually a modified web server it is sensible to enforce some security measures.



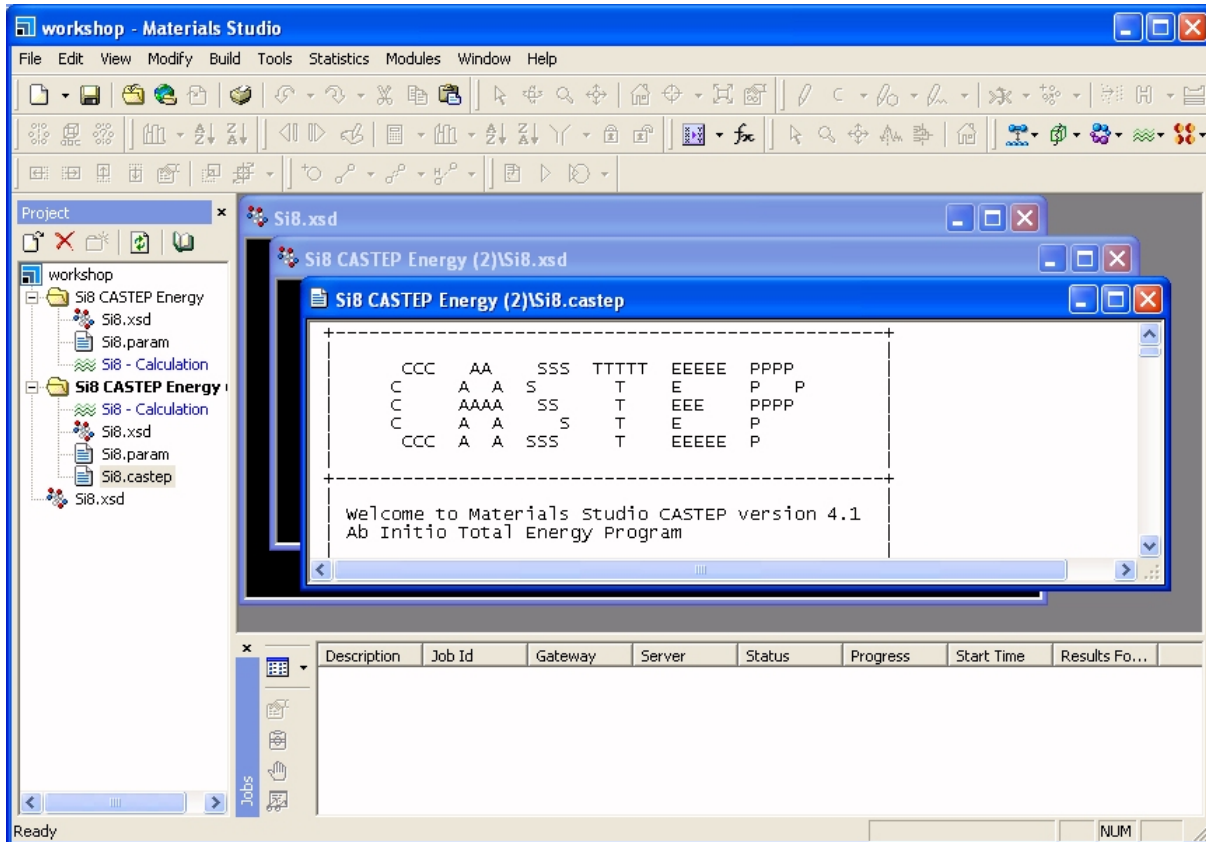
If your Gateway is password-protected (recommended), you'll need to enter your Gateway username and password (which are not necessarily the same as your Windows ones).



When the Castep job is running you will see its job ID and other details appear in the “job explorer” window. You can check its status from here, although our crude silicon calculation is so quick you probably won’t have time now.



Castep reports back when it is finished, and Materials Studio copies the results of the calculation back. The .castep file is opened automatically so you can see what happened in the calculation.



The main text output file from castep is displayed in Materials Studio. It starts with a welcome banner, then a summary of the parameters and cell that were used for the calculation.

After that, there is a summary of the electronic energy minimisation which shows the iterations Castep performed trying to find the groundstate density that was consistent with the Kohn-Sham potential. This is the so-called “self-consistent field” or “SCF” condition, and each line is tagged with “– SCF” so you can find them easily.

```

----- <-- SCF
SCF loop      Energy          Fermi          Energy gain    Timer    <-- SCF
                Energy          energy         per atom      (sec)    <-- SCF
----- <-- SCF
Initial      2.11973065E+002  4.85767974E+001          0.61 <-- SCF
Warning: There are no empty bands for at least one kpoint and spin; this may
slow the convergence and/or lead to an inaccurate groundstate.
If this warning persists, you should consider increasing nextra_bands
and/or reducing smearing_width in the param file.
Recommend using nextra_bands of 7 to 15.

  1  -7.22277610E+002  1.02240172E+001  1.16781334E+002  0.88 <-- SCF
  2  -8.53739673E+002  6.90687627E+000  1.64327579E+001  1.12 <-- SCF
  3  -8.62681938E+002  6.65069587E+000  1.11778315E+000  1.39 <-- SCF
  4  -8.62169156E+002  6.69758744E+000 -6.40977798E-002  1.72 <-- SCF
  5  -8.61880601E+002  6.78641872E+000 -3.60693332E-002  2.06 <-- SCF
  6  -8.61884687E+002  6.79549194E+000  5.10791707E-004  2.44 <-- SCF
  7  -8.61884645E+002  6.79874201E+000 -5.25062118E-006  2.75 <-- SCF
  8  -8.61884639E+002  6.79822409E+000 -8.40318139E-007  2.98 <-- SCF
----- <-- SCF

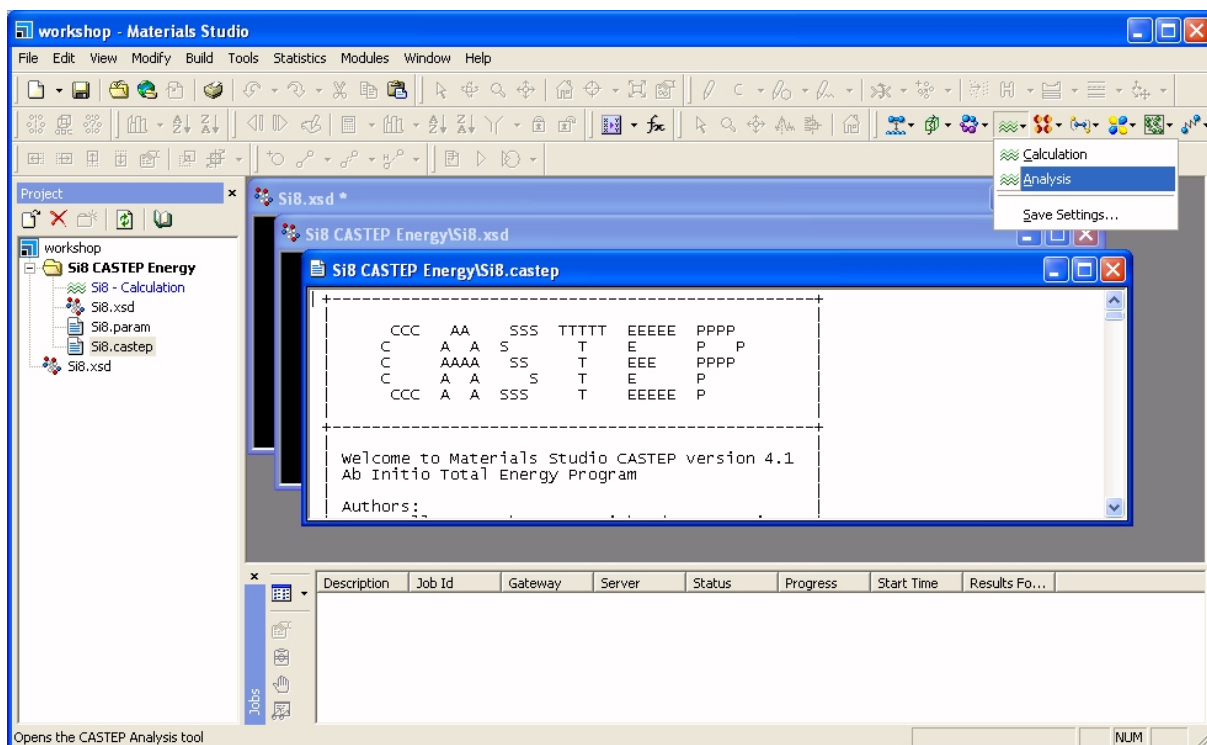
Final energy, E          = -861.8846385210    eV
Final free energy (E-TS) = -861.8846385210    eV
(energies not corrected for finite basis set)

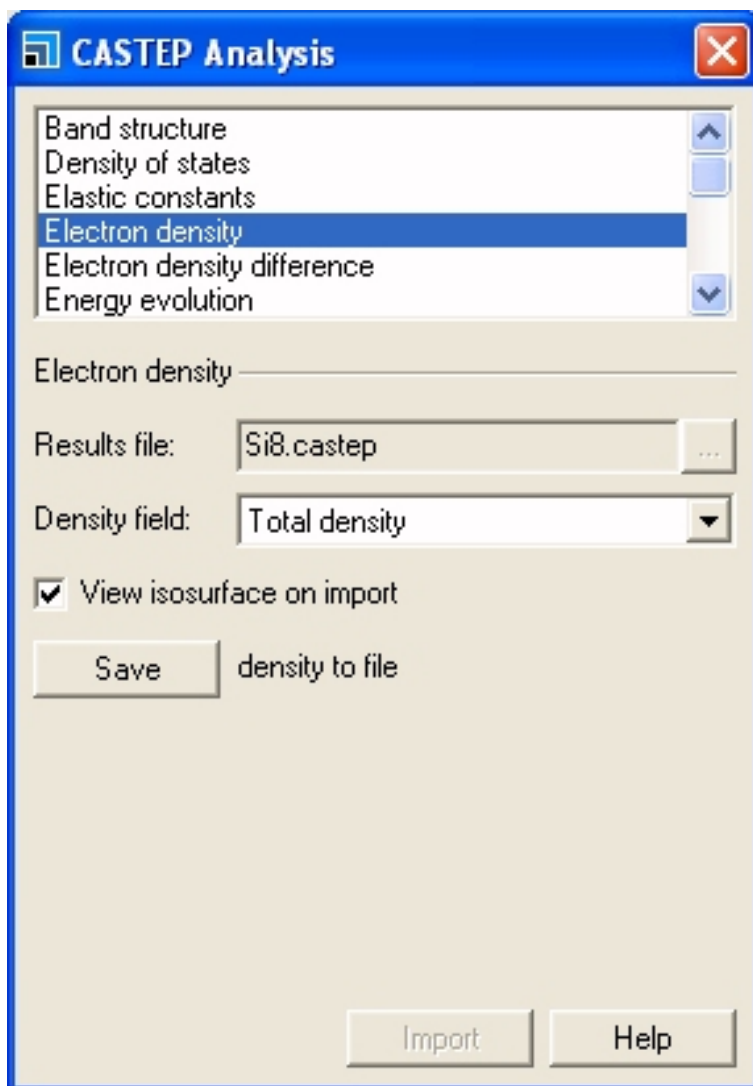
NB est. OK energy (E-0.5TS) = -861.8846385210    eV

```

We'll look at this output in more detail later. For now just note that the energy converges fairly rapidly to about 861.88eV, but that the energy is sometimes higher than this and sometimes lower.

Let's have a look at the calculated groundstate charge density.

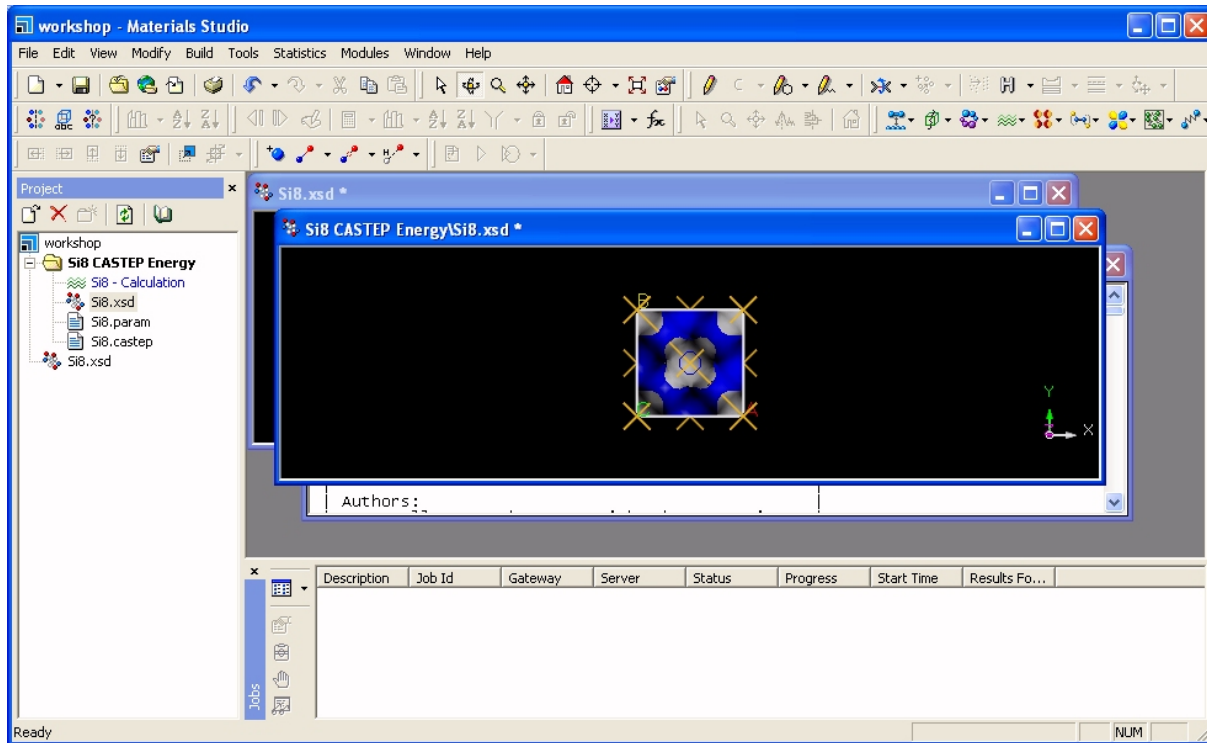




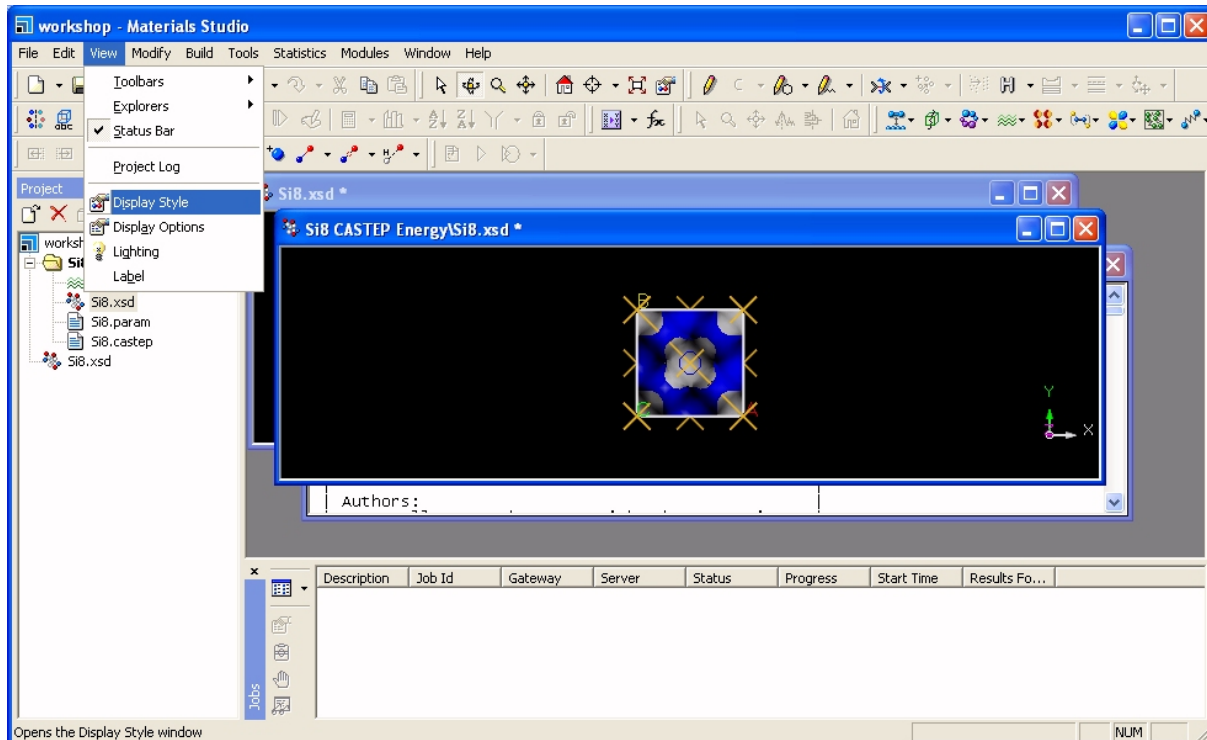
The Castep Analysis window lets you look at various properties you might have calculated during the Castep job. Select “Electron density”. Notice there’s a “Save” button which lets you write the density out to a text file so you can analyse it with another program. We don’t need this now, so just click on “Import”.

WARNING: amongst the properties listed here are “Band structure” and “Density of states”. If you select one of these from an energy calculation, Materials Studio will plot the band structure/DOS, but it takes the eigenvalues and k-points *from the SCF calculation*, not a proper band structure or DOS calculation.

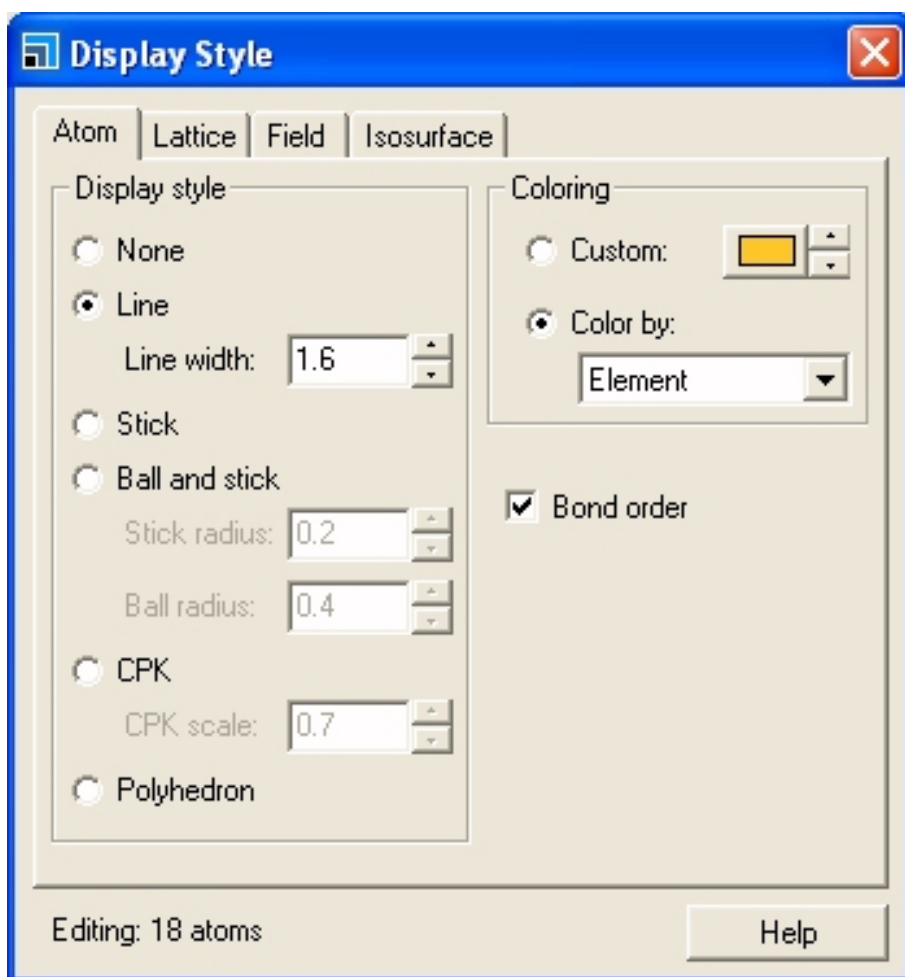
By default an isosurface of the charge density is overlaid on your simulation cell.

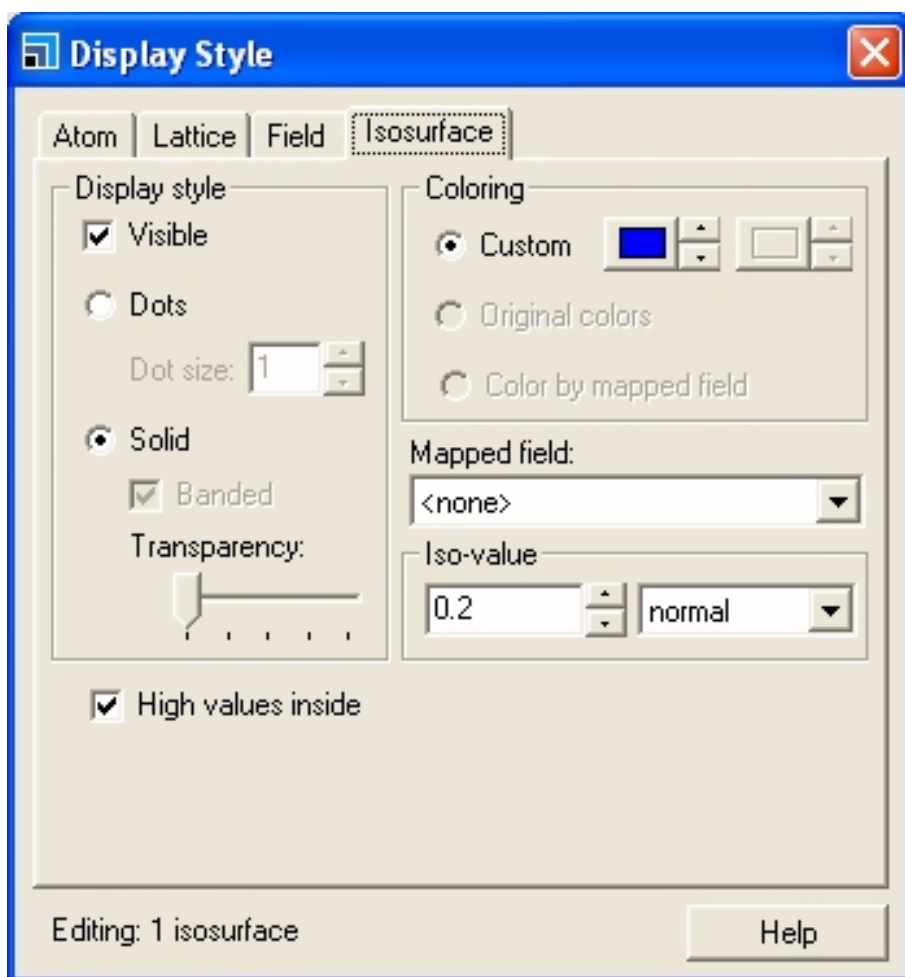


To change the isosurface Materials Studio is plotting, you need to change the “Display style”. Either use the right mouse button when the cursor is over the simulation cell, or use the drop-down menus:



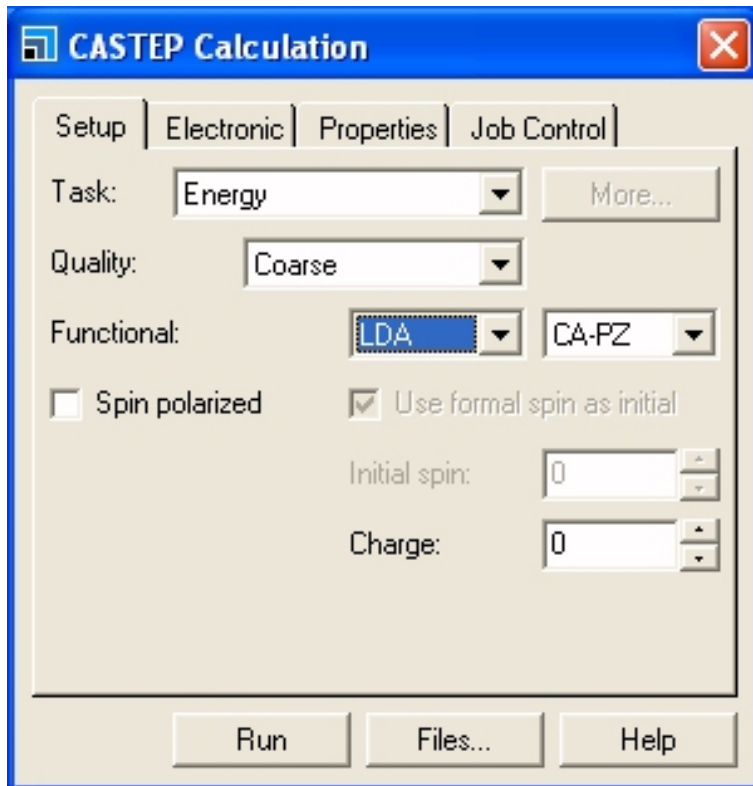
Notice that this is also the place you need to come to if you want to change the atom colouring or representation (e.g. from crosses and lines to ball-and-stick).



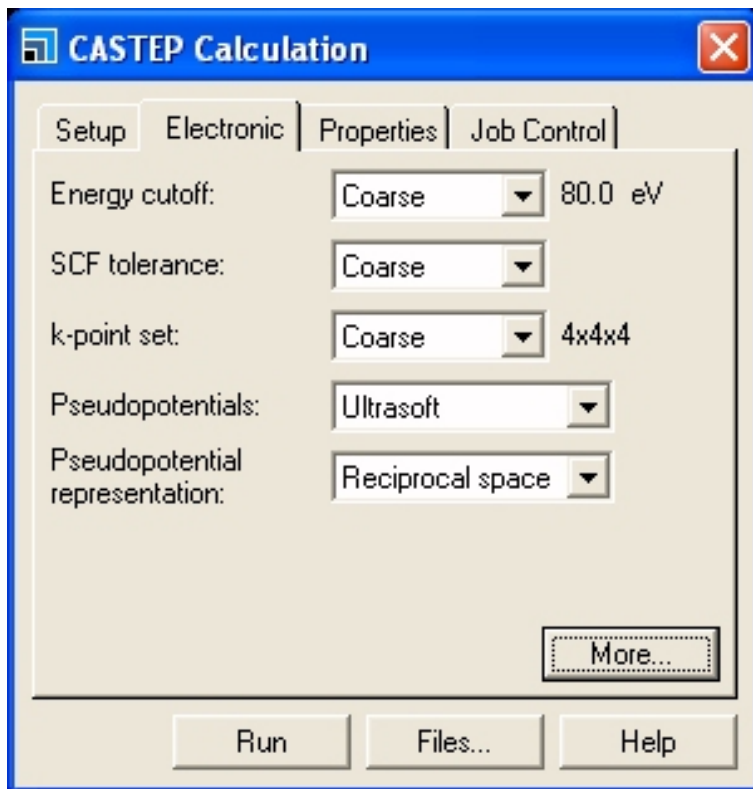


Try changing the value of the isosurface your plotting, to see where the charge density is greatest and least.

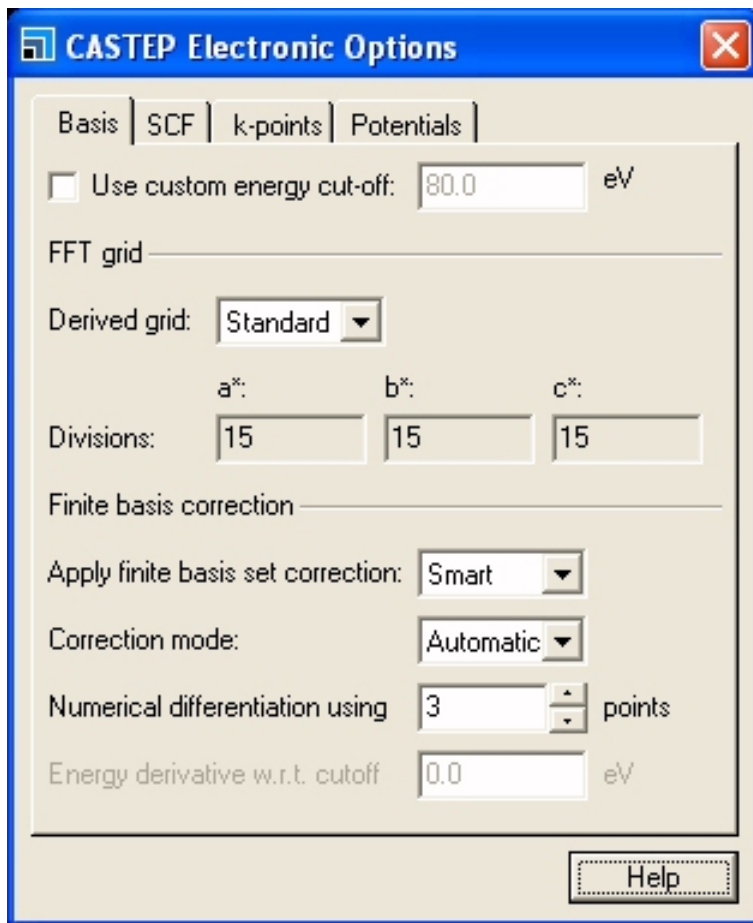
Hopefully you've now got the hang of the basic interface. Go back to your simulation system and open up the Castep window again.



This time select the “Electronic” tab.



This tab has a little more detail, and actually tells you what cut-off energy and k-point grid Castep will use for the given settings. Nevertheless we usually want finer control than this, so click on “More” .



Now at last we have four tabs that let us set some of the convergence parameters directly.

- Basis
Allows you to set a cut-off energy, as well as control the finite basis set correction.
- SCF
Sets the convergence tolerance for the ground-state electronic energy minimisation, as well as details of the algorithm used.
- k-points
Controls the Brillouin zone sampling directly. You can either specify a grid, or a desired separation between k-points.
- Potentials
Allows you to change the pseudopotentials used for the elements in your system.

In fact if you double-click on your param file in the project window you can edit it directly, but we'll restrict ourselves to using the GUI for now.

Before we continue, here's a quick recap of the basic approximations we use when performing practical DFT calculations:

- **Exchange-correlation (XC) Functional** - we don't know the exact density functional, so we have to approximate it. There are two common approximations:
 - **LDA** - the Local Density Approximation assumes the XC at any point is the same as that of a homogeneous electron gas with the same density.
 - **PBE** - this is a "Generalised Gradient Approximation" (GGA) and includes some of the effects of the gradient of the density.

You might think PBE is always better than LDA, but that's not true, both are approximations. You should try each one before deciding which is appropriate to your research project.

- **Basis set** - the wavefunction is represented by an expansion in a plane-wave basis. In theory the basis set required is infinite, but since the energy converges rapidly with basis set size we can safely truncate the expansion. The size of the basis set is controlled by the **cut-off energy**.
- **Brillouin zone sampling** - calculating the energy terms requires us to integrate quantities over the whole of the first Brillouin zone. In practice we approximate these integrals by sums over a discrete set of **k-points**.

Exercise 1. Using the Basis and k-points tabs, investigate how the calculated energy of the simulation cell converges with increased cut-off energy, and increased k-point sampling density. Why do they show these trends?

Exercise 2. Create a unit cell for bulk aluminium. Aluminium is also FCC, with space-group FM-3M and a lattice constant of about 4.05 Å. Investigate convergence of the calculated aluminium energy with respect to cut-off energy and k-point sampling.

Compare the total electronic energy with the total electronic free energy for both silicon and aluminium. Why do they differ for one and not the other?

During your calculations you might see a warning like this in the castep output:

```
Warning: There are no empty bands for at least one kpoint and spin; this may
slow the convergence and/or lead to an inaccurate groundstate.
If this warning persists, you should consider increasing nextra_bands
and/or reducing smearing_width in the param file.
Recommend using nextra_bands of 7 to 15.
```

Recall that the electronic energy minimisation algorithms need to include the entire set of occupied states. If the highest state you've included in the calculation is occupied, Castep has no way of knowing whether the next state should also have been occupied, and so recommends you include more bands. Only when the highest state is unoccupied can Castep be sure that all of the occupied bands have been included.

You can change the number of “empty” bands included in the Castep calculation from the SCF tab of the Castep Electronic Options window of Materials Studio, or just by editing the param file directly.

Exercise 3. Repeat the energy convergence test with respect to k-point sampling for aluminium, but using a smearing of 0.5eV (see the SCF tab; the default is 0.1eV). Feel free to use either Materials Studio, or direct editing of the param and cell files. You will probably need to increase the number of empty bands to 8 or so.

Compare the results with the previous aluminium calculations. Why the difference? Choose a particular k-point sampling density and look at the final total energy, free energy, and estimated zero temperature energy for the 0.5 eV smearing and compare them to the results with the original smearing.

Exercise 4. Go back to your silicon calculation, and look at the SCF tab on the CASTEP Electronic Options window. We're using the "Density Mixing" algorithm, and if you click on "More" you'll see we're using a Pulay mixing scheme with a charge mixing amplitude of 0.5. Investigate what happens as you vary this initial amplitude from close to 0 to close to 1.

The Pulay algorithm takes over after the first few SCF cycles, and overrides the mixing charge amplitude. This is not true of the Kerker scheme. Use the "More" button and change the mixing scheme to Kerker, and investigate the effects of the mixing charge amplitude again.

Exercise 5. Have a play with the Castep interface and Castep. Why don't you see whether you can get Castep to fail to converge? Remember what causes density mixing to be unstable: metals, degeneracies (band-crossings), multiple spin states, long cells, small smearing

widths etc. The only restriction is computational time, so if you make a large cell try not to have too many atoms in it or Castep won't finish in time!

If you manage to make Castep fail to converge, try to fix it by varying the DM parameters. If that doesn't work, how does EDFT do? Remember you can always save your cell and param files and copy them to Lagavulin if your PC isn't fast enough.

You can also try Castep on your favourite system. Things you might find useful:

- Materials Studio ships with lots of sample structures, just click “File” then “Import” and have a look, or create your own.
- To create a supercell from a unit cell, click on the “Build” menu, then select “Symmetry” and then “Supercell” .
- To modify atoms just left-click (or drag-select) to select them, and then you can use the “Modify” menu to change their element.
- Materials Studio has a useful surface builder so you can cleave crystals along bizarre planes without too much effort.

If you're stuck for things to do:

- Try making a supercell of two aluminium FCC cells, and swapping one of the aluminium atoms for erbium. Run that, and see what happens. Can you improve it?
- Use the task “properties” in the Castep window to calculate the DOS and band structures of silicon and aluminium. Now create a simple molecule surrounded by vacuum, and calculate its band structure and DOS. Do you get what you expect?
- Calculate the binding energy of a simple molecule. Run Castep for the molecule, and then again for a single atom of each of the elements in turn. Subtract the energies, and see what you get. How does the result change if you change the cut-off energy for (a) one calculation; (b) all the calculations?