

Composition spaces and phases in THERMOCALC

Eleanor C.R. Green (eleanor.green@unimelb.edu.au)

February 2020

In order to get meaningful results out of THERMOCALC – or maybe to get any results at all – it's crucial to understand how THERMOCALC handles chemical composition. This handy guide is possibly the most important piece of THERMOCALC documentation available. In particular, it explains how THERMOCALC defines composition within phases, what order variables do in solid solutions, and how to set the bulk composition and interpret the modal proportions of minerals and fluids correctly. It also covers information more broadly related to composition spaces, such as setting phases "in excess", obtaining the results of calculations with negative amounts of a mineral phase, and how to calculate solvi. It's written for practical use with a minimum of technical information, and will be updated in response to user queries and developments in the program.

Contents

1 Quick-reference: scripts mentioned in this document	2
2 Bulk composition	3
2.1 Units and normalisation	3
2.2 Scripts	3
2.3 Ferrous/ferric iron	4
2.4 T - X and P - X pseudosections	4
2.5 T - M_O pseudosections	4
3 Modal proportions ("modes") of phases	5
3.1 Meaning	5
3.2 Units	5
3.3 Negative modes and the <code>c8</code> script	5
4 Composition and order-disorder in phases	6
4.1 THERMOCALC's internal compositional variables	6
4.2 Order variables	7
4.3 Setting compositional and order starting guesses	7
4.3.1 In phase diagram calculations	7
4.3.2 In multiple-reaction thermobarometry	8

4.3.3	“Illegal Q starting guess” errors	8
4.3.4	Multiple Q solutions – ilmenite and omphacite	9
4.3.5	The range script	9
4.4	Output of molar oxides, site fractions and melt wt %	9
4.5	Working with fewer chemical components (“boiling”)	10
5	Phases co-existing across solvi	11
6	Always-present (“in excess”) phases	12
7	Corrigenda	12

1 Quick-reference: scripts mentioned in this document

script	section
% (comment)	2.2
c8	3.3
convliq	4.4
range	4.3.5
rbi	2.2
samecoding	5
setbulk	2.2
setexcess	6
tozero	6
which	6
xyz	4.3.2
xyzguess	4.3.1

2 Bulk composition

2.1 Units and normalisation

For many phase diagram calculations, the user must set one bulk composition. Bulk compositions are given in *molar oxides*. THERMOCALC will automatically normalise values to 1 or 100 % wherever necessary.

2.2 Scripts

Bulk compositions are usually set using the `setbulk` script, e.g.:

```
setbulk yes 5.44 50.09 8.69 11.67 12.15 7.78 2.50 0.47
```

Values are given in mole % oxides. The oxides must be *the same* as the oxides present in the *x-eos* used for the calculation, and they must be specified in the *order* that THERMOCALC expects (currently H₂O, SiO₂, Al₂O₃, CaO, MgO, FeO, K₂O, Na₂O, TiO₂, MnO, O, Cr₂O₃ for common rock-forming minerals, although you will not usually use all of them). If you aren't sure which oxides you need, or in which order, run the calculation without a bulk composition set and THERMOCALC will tell you.

The user can (should?) add comments, which will not be read by THERMOCALC, into the scriptfile to clarify which oxides are present and what the bulk composition represents:

```
% -----  
% sample 104-16  
% -----  
% molar oxides: H2O SiO2 Al2O3 CaO MgO FeO Na2O O  
setbulk yes 5.44 50.09 8.69 11.67 12.15 7.78 2.50 0.47  
% FeO: all iron as FeO  
% O: Fe2O3 (if Fe is the only cation with variable valence)  
% -----
```

Comments in THERMOCALC start with % and continue until the end of the line.

An alternative to the `setbulk` script is the `rbi` script block, which allows the user to specify a bulk composition in terms of the compositions and proportions of phases. Use of the `rbi` script block is shown in Richard White's tutorial 7, downloadable from the Pseudosections webpage, although note that the format has changed slightly in THERMOCALC 3.50.

2.3 Ferrous/ferric iron

Most model rock systems contain iron as the only multivalent element, with both ferrous and ferric iron present. The relevant variables in a THERMOCALC bulk composition are then FeOt (all iron as FeO) and O (oxygen). Via the reaction $2 \text{FeO} + \text{O} = \text{Fe}_2\text{O}_3$ we can see that the moles of iron oxides are given by

moles iron oxides	oxides in setbulk
FeO	FeOt - 2 O
Fe ₂ O ₃	O

2.4 T - X and P - X pseudosections

To calculate a T - X or P - X pseudosection, the user sets two setbulk scripts, e.g.

```
% -----  
% molar oxides:  H2O  SiO2  Al2O3  CaO  MgO  FeOt  Na2O  O  
  setbulk yes    5.44 50.09 8.69  11.67 12.15 7.78 2.50 0.47 % bulk 1  
  setbulk yes    1.68 48.96 9.90  13.82 16.56 6.85 1.67 0.30 % bulk 2  
    % FeOt: all iron as FeO  
    % O:    Fe2O3 (if Fe is the only cation with variable valence)  
% -----
```

THERMOCALC then allows the user to do calculations along the compositional vector that connects the two bulk compositions. It gives X -axis output in the form of a normalised distance between the two bulk compositions, with a scale of 0–1. This procedure is described in the pseudosection tutorials.

2.5 T - M_O pseudosections

A common type of T - X pseudosection is a T - M_O plot. This is calculated to find out how predicted phase relations depend on the assumed ferrous-ferric ratio. P - M_O pseudosections can also be calculated, of course.

For a T - M_O plot, the setbulk script block might look like:

```
% =====  
%          ... FeOt  ... O  
setbulk  ...  10    ...  1  % X_Fe3+ = 0.2  
setbulk  ...  10    ...  3  % X_Fe3+ = 0.6  
% =====
```

with all the other oxides kept constant.

As noted above, THERMOCALC gives X -axis output in the form of a normalised distance between the two bulk compositions, with a scale of 0–1. However, for T – M_O pseudosections (or other examples in which only one oxide is changing, e.g. T – M_{H_2O} pseudosections), it's usually more helpful to redraw the axis as M_O , where M_O is simply the value of O used in the calculation.

Alternatively (or as well), a more geologically intuitive variable might be:

$$X_{Fe^{3+}} = \frac{Fe^{3+}}{Fe^{2+} + Fe^{3+}} = \frac{2 Fe_2O_3}{FeOt} = \frac{2 O}{FeOt}. \quad (1)$$

3 Modal proportions (“modes”) of phases

3.1 Meaning

When the calculation involves a specified bulk composition, the modal proportions or “modes” of phases are calculated, in addition to the phase compositions. The modes embody the requirement of mass balance.

3.2 Units

Modes of phases are given as mole fractions on a 1-atom basis. This is loosely equivalent to a volume fraction, insofar as it makes it possible to compare e.g. garnet with a 20-atom formula unit with quartz on a 3-atom formula unit. However it's not strictly equivalent, since it doesn't account for the relative densities of crystallographic structures.

3.3 Negative modes and the `c8` script

Sometimes, THERMOCALC is able to calculate an equilibrium among the phases that the user asks for, but one or more of the resulting modes is negative. The calculation is therefore unphysical. By default, THERMOCALC will not print the solution, but will give the error code '8'.

However, sometimes the user is interested in the results of this calculation, particularly if they are struggling to find successful calculations, and would like to find a new `xyzguess` script block (section 4.3.1) to use in another attempt. You can make THERMOCALC print output from calculations with negative modes, by setting the script `c8 yes`. Of course, it's important to be careful with this script, and avoid plotting pseudosection boundaries involving negative modes. THERMOCALC prints a warning symbol, probably '#', on-screen next to calculations that it's worried about, which includes calculations with negative modes.

4 Composition and order-disorder in phases

4.1 THERMOCALC's internal compositional variables

Internally, THERMOCALC works with a set of variables for each x -eos (model phase) that is defined by the developer of that x -eos. These are commonly given single-letter names such as x , y , j , and their values are printed in the console, `-o` and `log` files when THERMOCALC is run, e.g.

```
<=====>
phases: pl ol opx cpx
-----
P (kbar)      T (C)      ca (pl)      k (pl)      x (ol)      c (ol)      Q (ol)
  5.0000    1100.000    0.7162    0.003451    0.1021    0.003846    0.0002174
              f (opx)      t (opx)      cr (opx)      j (opx)      x (cpx)
              0.02740    0.01538    0.03555    0.005381    0.09109
```

Each variable is defined in terms of one or more site fractions. In THERMOCALC 3.50 and above, the definitions are given in the `-it` output file, and they can also be found in the original papers. In all but the most complex x -eos, the variables can be straightforwardly divided into *compositional* (sometimes called *macroscopic*) variables and *order* variables. Compositional variables are related to mineral analytical composition, and where possible we try to define them in a way that seems geologically natural – for example, in most x -eos, x will be defined as total FeO/(FeO+MgO) for the phase. Order variables (usually beginning with Q in a solid or y in a fluid) then define the details of the partitioning of ions across sites, as discussed in 4.2. The set of variables is chosen so that the variables fully and uniquely specify the site fractions of each of the mixing ions. The algebraic expressions for the site fractions, written in terms of the THERMOCALC variables as shown in the `-it` output file, are innately normalised to a constant number of ions/vacancies per site (for crystalline phases), and maintain net charge across mixing sites.

The set of variables that a developer chooses for an x -eos is non-unique. We generally try to use the same set of variables for any given x -eos wherever it appears. However, we do sometimes change our decisions over time. In the case of isosymmetric plagioclase and K-feldspar, we use different variables to characterise what is otherwise the same x -eos, to facilitate reducing the x -eos into fewer components (discussed in 4.5).

Although the THERMOCALC variables aren't always geologically intuitive, they are the values that are printed on-screen (for a number of reasons). Therefore they are the easiest way for the user to keep an eye on the phase compositions during calculations. If phase compositions drift into an unlikely part of composition space, it can indicate that there is a more stable assemblage that the user has forgotten to look for. When this happens, the compositions can eventually run into chemical subsystems, causing THERMOCALC to stop calculating.

4.2 Order variables

In many *x*-eos, pairs of elements are able to mix on more than one site, most commonly (Mg,Fe) mixing on, for example, both M1 and M2 in the pyroxenes. THERMOCALC then needs additional variables, in order to fully specify how the elements are arranged on the mixing sites, for a known composition. These are called order variables. They usually start with *Q* in crystalline phases or *y* in fluids, and they are identified in the *x*-eos description in the *-it* file.

Generally, order variables are not of much interest to the user, and can be ignored. However, they're sometimes relevant if the user is struggling with "illegal" starting guesses, see [4.3.3](#).

4.3 Setting compositional and order starting guesses

4.3.1 In phase diagram calculations

In phase diagram calculations, THERMOCALC solves a set of equations to find the equilibrium compositions and states of order of the phases in the user-specified assemblage. The equations are non-linear, so it must solve them iteratively, and it must start from some realistic values for the compositional and order variables. These are called "starting guesses", and the user feeds them to THERMOCALC via the `xyzguess` script block, which specifies values for each variable in each *x*-eos in the desired assemblage:

```
% -----  
xyzguess x(opx)          0.539535  
xyzguess y(opx)          0.0502431  
xyzguess f(opx)          0.0226851  
xyzguess c(opx)          0.0486599  
xyzguess Q(opx)          0.387159  
% -----  
xyzguess x(g)            0.790028  
xyzguess z(g)            0.354536  
xyzguess f(g)            0.0339621  
% -----  
xyzguess x(ol)           0.15  
% -----  
:
```

Starting guesses are one of the more difficult parts of the THERMOCALC user experience, but it's crucial to be aware of the problem. If the starting guesses are not good enough, THERMOCALC may either fail to find any equilibrium when one exists, or may find an equilibrium that is not the most stable one

present.

To help users navigate the problem, we provide sets of starting guesses for the *x-eos* (that *might* be good enough) in the thermodynamic input file distributions available from the *HPx-eos and THERMOCALC* website. Thereafter, THERMOCALC prints the results of successful calculations in the log file in *xyzguess* form, and this block can be copy-pasted into the scriptfile as a starting point for subsequent calculations. If the scriptfile contains multiple *xyzguess* script blocks, THERMOCALC uses the first block for each phase. If it complains that there are too many *xyzguess* scripts, an asterisk, *, can be added to the scriptfile to make THERMOCALC stop reading the file below that point.

The user must avoid setting starting guesses that are “illegal”, resulting in non-physical site fractions that lie outside the range 0-1: for example, setting $(z(hb) + c(hb) = x_{Na}^{M4} + x_{Ca}^{M4}) > 1$. See the *-ic* file, or the *x-eos* descriptions provided with the thermodynamic input file distributions, for definitions of the variables and expressions for the site fractions. Section 4.3.3 discusses cases of illegal starting guesses in which the compositional variables alone do not invalidate the site fractions, but the combination of compositional and order variables does. Note that any *xyzguess* script block copied out of the log file will always be legal.

4.3.2 In multiple-reaction thermobarometry

A script block similar to the *xyzguess* script block is also needed in multiple-reaction thermobarometry. Here, however, it serves a different purpose, and to denote this, the keyword *xyzguess* is replaced with simply *xyz* (in fact, if you forget, THERMOCALC will happily use *xyzguess*). Here, however, the compositional variables supplied by the user are not starting guesses; they are the analytical mineral compositions in the rock sample. The analytical compositions must be translated into THERMOCALC variables, first converting into cations and assigning to sites to obtain the site fractions.

4.3.3 “Illegal *Q* starting guess” errors

If the user inadvertently provides a set of starting guesses that would result in non-physical site fractions, THERMOCALC looks to see if there are any order variables involved. If there are order variables present, and if the compositional variables alone do not invalidate the site fractions (section 4.3.1), it should be possible to find starting guesses for the order parameters that make the *xyzguess* script block legal again. So THERMOCALC makes a few trials using random values for the order variables, looking for a set that produce valid site fractions.

If it finds such a set, it may go on to produce a successful calculation. However, the user should look out for messages complaining about “illegal *Q* starting guess for <phase>”, and should change the starting guesses to a legal set rather than allowing the situation to continue. If THERMOCALC continues to look randomly for order-variable starting guesses, then, on re-running the same calculation, it may not find a

legal set of Q starting guesses. Worse, it might use Q starting guesses that lead to a *different solution to the calculation* – see 4.3.4.

Notes:

- The ranges of Q -values that constitute legal starting guesses are a function of the compositional variables.
- If there are Q values present, THERMOCALC assumes that these are at fault. You may need to check, however, that the problem doesn't lie in the compositional variables, as in section 4.3.1.

4.3.4 Multiple Q solutions – ilmenite and omphacite

In the x -eos for ilmenite and the sodic-calcic pyroxenes dio, o and jd (diopside, omphacite, jadeite), it's possible to generate two solutions involving different values of Q , i.e. different degrees of order. We usually want the more stable solution, i.e. the one with the lower value of the Gibbs energy, G , which is printed on-screen. The user should identify the more stable value in a trial pair of calculations (as far as I'm aware, the more stable calculation is always associated with the larger value of Q), then use a `xyzguess` script block taken from the log file of the more stable calculation, so that later calculations are correct.

4.3.5 The range script

In the `xyzguess` script blocks, some variables may be followed by the script `range <min> <max>`. The `range` script specifies what range of values the variable is allowed to take inside THERMOCALC, with a default range of 0–1. The range may be larger than this – for example, some order variables are defined to have a range of $-1 < Q < 1$, where $Q = 1$ in the fully ordered phase and $Q = -1$ in the fully anti-ordered phase.

In general the user should leave the `range` scripts alone. It's certainly not advisable to reduce the specified ranges in order to stop variables taking on geologically-implausible values. This might make THERMOCALC stop calculating, but does not fix the problem, which indicates either that the user has missed out a phase from the “most stable” assemblage, or bad behaviour from the x -eos. However, occasionally the x -eos developer might have accidentally set the `range` variable too small, and arbitrarily prevented the x -eos from exploring the whole of its composition–order space. If you suspect this, please ask about it.

4.4 Output of molar oxides, site fractions and melt wt %

In THERMOCALC 3.46 onwards, all phase diagram calculations generate an output file with the suffix `-ic`. This contains phase compositions in the form of molar oxides and site fractions.

For melt phases only (*x*-eos called “liq”) it’s also possible to generate on-screen and logfile output in wt %. This is done using the following rather clunky script block, shown here for liquids in the NCKFMASHTO system:

```
%
      H2O SiO2 Al2O3 CaO  MgO  FeO  K2O  Na2O TiO2   O  Cr2O3
convliq  1   0   0   0   0   0   0   0   0   0   0   % H2O
convliq  0   1   0   0   0   0   0   0   0   0   0   % SiO2
convliq  0   0   1   0   0   0   0   0   0   0   0   % Al2O3
convliq  0   0   0   1   0   0   0   0   0   0   0   % CaO
convliq  0   0   0   0   1   0   0   0   0   0   0   % MgO
convliq  0   0   0   0   0   1   0   0   0   0   0   % FeO
convliq  0   0   0   0   0   0   1   0   0   0   0   % K2O
convliq  0   0   0   0   0   0   0   1   0   0   0   % Na2O
convliq  0   0   0   0   0   0   0   0   1   0   0   % TiO2
convliq  0   0   0   0   0   2   0   0   0   1   0   % O
convliq  0   0   0   0   0   0   0   0   0   0   1   % Cr2O3
```

For melts in smaller systems, remove the appropriate rows and columns.

4.5 Working with fewer chemical components (“boiling”)

It’s often useful to work with a smaller set of chemical components than the maximum number found in the *x*-eos. We do provide axfiles in a variety of chemical systems, but what happens if the user wants to work in a system for which there is no axfile? It is sometimes possible to eliminate chemical components from an *x*-eos, in a process that for some reason we call “boiling out” components.

To do this, start with the definitions of the compositional variables of all of the *x*-eos you want to use, as found in the `-it` file or THERMOCALC thermodynamic input file distribution. It is possible to boil out a component – say, Cr_2O_3 – from an *x*-eos if (1) there are compositional variables involving Cr_2O_3 , and (2) by setting them to zero, you would set the proportions of all Cr-bearing end-members zero. This is true for all of the Cr-bearing *x*-eos in the igneous family of HP*x*-eos, so you could boil out Cr_2O_3 by setting *ek(L)*, *cr(g)*, *cr(opx)*, *cr(cpx)* and *c(spn)* to zero. This is done in the `xyzguess` script block:

```
xyzguess x(opx)          0.25
xyzguess y(opx)          0.05
xyzguess c(opx)          0.05
xyzguess Q(opx)         -0.04
xyzguess f(opx)          0.02
xyzguess t(opx)          0.01
xyzguess cr(opx)         0      % boiler
xyzguess j(opx)          0.01
```

Multiple compositional and order variables may need to be set to zero in one x -eos – for example, boiling out FeOt usually requires the elimination of variables related to Fe-Mg order-disorder. Take care!

It's not necessarily possible to boil the x -eos as you would like. This may be because of the developer's choice of variables – for example, K or Ca can be boiled out of ternary $\bar{C}1$ plagioclase ($ca = x_{Ca}$, $k = x_K$), but Na cannot (in an emergency you could ask us to re-cast the x -eos into new variables). In other cases, it simply isn't possible to eliminate some components independently of other components.

5 Phases co-existing across solvi

Solvi are miscibility gaps between phases with identical symmetry. At high temperature, the solvus closes, leaving a single phase with continuously varying composition. Some well-known solvi include those between the clinoamphiboles hornblende, glaucophane and actinolite, the clinopyroxenes diopside, omphacite and jadeite, as well as the peristerite gap which is discussed at the end of this section.

Since at high temperatures there is only one phase present, solvi need to be modelled using a single x -eos for the phase in question. To do this in THERMOCALC, use the `samecoding` script to replicate the relevant x -eos. For example, the script

```
samecoding hb act gl
```

replicates the hb (hornblende) x -eos twice, assigning the names gl and act to the copies. Both hb and the new phases created, gl and act (representing glaucophane and actinolite), must be given their own sets of starting guesses, just like any other phase. For gl and act, start by copying the `xyzguess` block for hb and changing the phase name, then adjust the starting guess values to reflect the correct composition, for example setting $z(gl) \gg z(act)$, where $z(\text{clinoamphibole}) = x_{Na}^{M4}$.

When doing calculations with phases that have solvi, it's particularly important to keep an eye on the calculated compositions, and make sure they are on the side of the solvus that you want. If the calculation involves two phases separated by a solvus, check that THERMOCALC is actually calculating different compositions for these phases. If it is calculating essentially the same compositions for both phases (it may try to alert you by printing '#' next to the calculation), either the starting guesses for one of the phases must be corrected manually, or you may now be at a temperature above the solvus, in which case you should remove one of the phases from the assemblage.

N.B. the x -eos currently used for the feldspars don't recreate the peristerite gap. Instead, it seems surprisingly successful to use the usual metamorphic $\bar{C}1$ plagioclase x -eos, combined with the pure albite end-member, ab.

6 Always-present (“in excess”) phases

Sometimes, in phase diagram calculations, the user knows that some phases will appear in most or all of the fields, e.g. quartz, muscovite and water (q, mu, H₂O) across much of the subsolidus for a metapelite. If the script `setexcess q mu H2O` is set, THERMOCALC will always include these phases in a calculation; the user will not have to specify them in the `which` script or in interactive output.

If a fluid phase such as H₂O is in excess, the user might want to see the modes of the mineral phases only, with the fluid normalised out. This can be achieved by setting the script `tozero H2O`.

Note: prior to THERMOCALC 3.50, it was possible to omit H₂O from the bulk composition and apply `setexcess H2O`, with THERMOCALC printing only the output in which H₂O is normalised out.

Internally, THERMOCALC would add a large amount of H₂O, e.g. 100 moles, to the bulk composition before doing the calculation. This facility has been removed, because we are trying to avoid THERMOCALC doing things behind the scenes that the user is not aware of.

7 Corrigenda

Corrections to previous versions of this document:

None yet...