

Urgently seeking equilibrium: I. Rapid and robust phase saturation calculations with the Metastable EXchange eQuilibrium ALgorithm (MEXQAL)

Aaron S. Wolf^{*1} and Mark S. Ghiorso²

¹University of Michigan, Ann Arbor

²OFM Research

* Correspondence: Aaron S. Wolf <aswolf@umich.edu>

Abstract

A nice abstract.

Alternate possible titles:

- Find equilibrium fast every time
- A rapid and robust search for equilibrium

Significance or key points

- Key point 1
- Key point 2
- Key point 3

Keywords: thermodynamics, modeling, algorithm, equilibrate

1 Introduction

Key Points

- see Appendix section 6.3 for exhaustive list to be paired down/simplified

[gist: Fast & accurate equilibrium critical to thermodyn simulation] Reliable determination of equilibrium is critical for thermodynamic simulations. Predicting the evolving state of the system hinges entirely on the speed and accuracy of the underlying algorithms responsible for finding the equilibrium states of each phase. Given the large abundance and variety of phases relevant to geological and planetary simulations, this demands a general technique that is sure to rapidly converge to the optimal answer, regardless of details unique to specific phase.

[gist: MEXQAL establishes equilibrium properties of individual phases in an assemblage; necessary for finding equilibrium] **[NOTE-ASW:Big claim needs to be demonstrated!]** We develop the Metastable EXchange eQuilibrium ALgorithm (MEXQAL), which calculates saturation state phase properties for any thermodynamic solution model and outperforms all existing published algorithms. Ghiorso (2013) presented a generalized saturation state algorithm in which metastable exchange equilibrium was considered in the context of an omni-component phase (like melt or aqueous fluid), which can freely exchange atoms with all other coexisting (or precipitating) phases in the assemblage. Building upon this previous work, we seek to broaden the intended scope for this algorithm to the general case that lacks an omni-component phase. Absent an omni-component phase, the stable assemblage consists of a collection of phases that together minimize the total Gibbs energy while maintaining the total bulk composition of the system. Even in this case, a metastable exchange algorithm is highly useful, as briefly discussed in Ghiorso (2013) which introduced the concept of a hypothetical omni-component phase **[NOTE-ASW: say why/how it's useful!]**. Here we consider arbitrary phase assemblages which can generally lack an omni-component phase, where a metastable exchange algorithm can allow us to determine the most energetically favorable composition for each solution phase given a guess of the elemental chemical potentials of the system. With such a computational tool, we can then iteratively determine the elemental chemical potentials by minimizing the total Gibbs energy of the system.

[NOTE-ASW: Improve topic sentence.] [*gist: primary applications, subsolidus assemblage & miscibility gaps*] There are numerous standard thermodynamic calculations that could make immediate use of a computationally efficient algorithm of this type. Within geological applications, calculating subsolidus phase assemblages is a typical application, used to determine in situ chemical and physical properties of a rock and extreme T/P conditions. Similarly, we can also use this algorithm to determine the equilibrium assemblage for a single solution phase that undergoes exsolution, where the bulk composition potentially lies inside of a miscibility gap. This method would be a critical part of the overall algorithm enabling the efficient and accurate calculation of tielines and 3-phase triangles for an immiscible bulk composition (as well as potential higher order coexisting assemblages). Metastable exchange equilibrium thus provides a useful reference frame for modeling equilibrium processes, even in the absence of an omni-component phase. [NOTE-ASW: Use affinity to use phase present/absent constraints for calibrating rxn boundaries like solidus position, olivine in, etc]

[NOTE-ASW: Also useful for evaluating quality of an inferred geologic simulation, based on assemblage comparison, undersaturated affinity acts as penalty weighting factor.]

2 Methods

[NOTE-ASW: Old title # Metastable EXchange eQuilibrium ALgorithm overview] [*gist: Intro to algorithm, it's benenfits and test applications*] Our new Metastable EXchange EQilibrium ALgorithm (MEXQAL) improves and generalizes the saturation state algorithm presented in Ghiorso (2013). It utilizes a similar iterative analytic approach, but simplifies the composition update calculation while utilizing additional derivative information in a rapid approximate optimization scheme. These changes improve numerical accuracy while cutting runtime by over an order of magnitude, reducing the number of Gibbs energy evaluations and remaining accurate for both major and trace components. The details of this algorithm are presented in this section. This section describes the details of this new algorithm, demonstrating the improvements using relevant thermodynamic models for geological phases.

2.1 Conditions of metastable exchange equilibrium

[NOTE-ASW: Overly long description of phase affinity, saturation, and metastable equilibrium; need to cut down] Phase saturation describes whether a phase is expected to be present in the equilibrium assemblage of a system at a particular set of conditions. Saturated phases appear in the assemblage because they help to minimize the relevant thermodynamic potential (e.g. total Gibbs energy at defined T, P, and bulk composition). Any equilibrium phase assemblage can be described in terms of a set of chemical potentials of the components of the system. Usually it is simplest to think in terms of elements (or oxides in many geological applications). This chemical potential vector describes a plane in composition-energy space where all of the saturated phases (present in the assemblage) lie somewhere along this plane. All metastable phases necessarily lie above this plane, since they are energetically disfavored, and no phases lie below the plane if the phase assemblage is indeed the equilibrium assemblage. The bulk composition of the system is likewise expressed as the weighted average composition of all phases present.

When performing thermodynamic modeling, especially for evolving systems or ones that may possess important kinetic barriers, it is critical to track not only the saturated phases, but the nearly stable undersaturated phases as well. The degree of metastability (or undersaturation) is described by the saturation affinity:

$$\mathcal{A} = -(\mu_i - \hat{\mu}_i)$$

where μ_i is the chemical potential vector for the phase in question, $\hat{\mu}_i$ is the external chemical potentials of the system, and the saturation affinity \mathcal{A} is the vertical offset between them. (Note that this expression holds for all components, all values of i .) The fact that the chemical potential plane of every phase lies parallel to—though possibly vertically offset from—the chemical potential plane of the system reflects that the phase is in metastable exchange equilibrium with the system as a whole [[202010231434]]. It is for this reason that every phase has a single saturation affinity that is component-independent. Additionally, the negative sign convention reflects the fact that under-saturated phases, which lie above the chemical potential plane of the system (with positive excess gibbs energy) are taken to have negative affinities indicating that they will not spontaneously appear in the system.

Using the standard convention, only reactions with positive affinities are allowed to progress forward and drive the system toward equilibrium.

[gist: quantify equilibrium condition for solution phase] We consider a solution phase in metastable exchange equilibrium with an externally imposed chemical potential ($\hat{\mu}_i$). For numerical stability, the imposed chemical potential is expressed as a deviation from the solution phase’s pure endmembers (μ_i^0):

$$\Delta\hat{\mu}_i = \hat{\mu}_i - \mu_i^0 \quad (1)$$

Similarly, the chemical potential of the solution phase is expressed as a deviation from ideal behavior:

$$\begin{aligned} \mu_i^{\text{xs}} &\equiv \mu_i - \mu_i^{\text{ideal}} \\ &= \mu_i - \mu_i^0 - RTm \log X_i \end{aligned} \quad (2)$$

where μ_i^{xs} is the excess non-ideal component of the chemical potential and RTm is the thermal energy scale modified by the atomic site multiplicity m . The equilibrium affinity condition for the solution phase is then given by:

$$\phi_i \equiv \Delta\hat{\mu}_i - \mu_i^{\text{xs}} = A + RTm \log X_i \quad (3)$$

where A is the saturation affinity of the phase relative to the imposed chemical potential, and $RTm \log X_i$ is the entropic ideal mixing contribution to the Gibbs energy. To obtain the equilibrium properties of the solution phase, we need to find the composition X_i that satisfies the above expression for all components, along with the closure condition that the composition is normalized ($\sum_i X_i = 1$). **[NOTE-ASW: Move this to section intro?]** This will simultaneously also provide the phase saturation affinity, indicating whether the phase is over-saturated ($A > 0$), under-saturated ($A < 0$), or simply stable ($A = 0$). Note that for an equilibrium assemblage calculation, we know that all phases present will be saturated (with affinities that are zero within tolerance), with all other possible phases having unfavorable negative affinities. **[NOTE-ASW: add caveat about dependent species]**

2.2 Species vs components in complex multi-site solutions

[gist: ideal mixing requires positive components] One of the critical nuances of the MEXQAL method (also discussed in Ghiorso (2013)) is that it requires all calculations be performed in dependent species space. This arises because the equilibrium affinity condition applies an ideal mixing correction to the saturation affinity to best match the estimated chemical potential offsets ϕ_i . The entropic mixing term has the familiar form of $RTm \log X_i$, implicitly assuming that the components are all positive and bounded between 0 and 1. This assumption is trivially satisfied for simple solutions, like the MELTS liquid phase or the Feldspars, because their solution endmembers are independent and fully span the physically realizable composition space.

[gist: project composition from endmember components to dependent species space] More complex solutions involving multi-site coupled-substitution, however, like for the clinopyroxenes or spinels, require the ability to describe compositions involving negative amounts of some components. These negative quantities are necessary to represent all possible species for the solution phase. A simple and straightforward fix for this problem is to simply project it into the larger dependent species space, rather than the original independent endmember component space. The full set of dependent species for any solution phase has the important property that any physically meaningful composition can be represented as a positive combination of the dependent species, reflecting the fact that they fully span the possible composition space. **[NOTE-ASW: Point out limitations of PerpleX approach, or at least reference to discussion section.]** By making this change, we guarantee that the fractions of each dependent species are always positive and we can make use of all of the same mathematical architecture developed for simpler solutions. This strategy for handling complex reciprocal solution phases is likewise discussed in the original Ghiorso (2013) paper, and we follow a similar procedure (for mathematical details, see Appendix section 6.1).

2.3 Direct iterative convergence

[NOTE-ASW: What is clearest terminology? Direct or progressive convergence? Direct update?]

[gist: Faster more accurate convergence w/ progressive method] While the above equation cannot be analytically solved directly for arbitrarily complex solution phases, it can be manipulated to allow a step-wise

analytic solution update; this enables an iterative approach to achieve a simultaneous self-consistent solution for composition and affinity. By rearranging and solving for the unknown phase composition, we obtain:

$$\begin{aligned} X_i &= \exp\left(\frac{\phi_i - A}{RTm}\right) \\ \hat{X}_i &\equiv \exp\left(\frac{\phi_i}{RTm}\right) \end{aligned} \tag{4}$$

where \hat{X}_i is the unnormalized compositional vector, with a constant of proportionality that depends on the unknown phase affinity. In this unnormalized expression, approximate values are known for every term since the current estimation of the chemical potential offset ϕ_i is readily calculated for the current composition estimate. The actual composition is then trivially determined by normalizing the compositional vector, directly producing updated estimates for both the phase composition and the saturation affinity:

$$\begin{aligned} X_i &= \hat{X}_i / \left(\sum_i \hat{X}_i\right) \\ A &= RTm \log\left(\sum_i \hat{X}_i\right) \end{aligned} \tag{5}$$

This approach yields mathematically equivalent results to the repeated fraction procedure described in Ghiorso (2013), while remaining simpler and improving numerical stability and accuracy. **(NOTE-ASW: Add appendix section describing repeated fraction method of Ghiorso 2013 for comparison.)** The self-consistent composition and affinity for the phase is determined using the above procedure to iteratively update the phase composition, and using that updated composition to determine an updated chemical potential offset ϕ_i . This is repeated until both the saturation affinity and the molar composition for each component converge to within the desired tolerance. We find that the compositional convergence shows optimal performance when monitored in log-space, ensuring that both major and minor or trace components converge to the same fractional degree (rather than absolute molar fraction). Although this iterative update scheme is robust, it can be further improved by making use of derivative information to dramatically speed up the calculation. **[NOTE-ASW: Do we need to quantify improvement over repeated fractions method?]**

2.4 Low-cost affinity updates via compositional-perturbation models

Chemical thermodynamic simulations, especially those involving compositional evolution or model calibration require many repeated calculations of the chemical potentials for each phase, typically dominating simulation computing time. Despite generally modest adjustments in composition from one iteration to the next, the computational burden for each of these calculations is fully realized in every function call, dramatically and unnecessarily increasing overall computing times. We are thus highly motivated to seek out simple local approximations that can achieve dramatically faster runtimes while retaining accuracy over reasonable ranges in composition. In this section, we present a novel local regular solution model which provides rapid and accurate phase energy estimates and compare it with the baseline ideal and linear solution models.

[NOTE-ASW: ‘linear’ term is confusing due to collisions with linear comp, log comp, and linearized eqn for least squares. Maybe ‘local expansion’ or ‘first order’ approximation rather than linear.]
[NOTE-ASW: similarly, ‘direct’ convergence method is confusing; perhaps ‘forward,’ ‘feed-forward,’ ‘progressive’ convergence instead?]

2.4.1 Basic first-order perturbation approximation

[gist: Gibbs curvature enables composition-dependent update of chemical potentials; but straightforward first order expansion is limited.] The progressive convergence method presented above adopts a conservative approach, where the current composition determines a chemical potential update, which in turn determines the new compositional update. In this procedure, the chemical potential evaluation is the computationally-limited step, especially for complex solution phases involving multiple ordering parameters (as is common for geologically important phases like the pyroxenes or spinels). We can speed up the process by also evaluating the Gibbs energy

curvature matrix, $\partial_n^2 G = d\mu_i/dn_j$, which expresses the (linear) compositional-dependence of the chemical potentials. This curvature matrix enables rapid approximate chemical potential updates, allowing us to get closer to the self-consistent answer without additional costly chemical potential updates. To achieve this, we adopt a nested convergence structure where approximate convergence is achieved within a fast inner loop, limiting the number of required iterations for the computationally intensive outer loop that updates chemical potentials and the curvature matrix. This proves to be a worthwhile optimization in typical applications because the additional computational cost of calculating the derivative matrix is outweighed by substantial iteration reduction.

A naive application of this information would simply impose a linear update of the chemical potentials after each compositional update. Though this does converge, the gains are rather modest since it approximates the Gibbs energy function with a quadratic surface, and completely misses the logarithmic nature of mixing entropy. This first-order expansion can be evaluated simply in molar composition space, or alternatively applied in log-composition space by transforming the curvature matrix (e.g., $d\mu_i/d \log n_j = n_j \cdot d\mu_i/dn_j$). Either way, this imposes a simple linear correction to the chemical potentials which improves performance, but is limited to a compositional region relatively close to the self-consistent final answer. Instead, we seek a method capable of confident extrapolation over large ranges in composition space.

2.4.2 Local regular solution approximation

[gist: Basic description of local regular solution.] [NOTE-ASW: Adjust topic sentence(s) to match new location. Overall paragraph very rough, needs reworking in new location.] *Regular solution models represent the next step in improving accuracy for chemical mixing models, and while they can be quite useful, their simplicity limits their applicability to complex ordered phases. Nevertheless, the ability of the regular solution to capture ideal mixing and roughly represent (in many situations) non-ideal excess interactions makes it attractive for calculations considering local perturbations.* We therefore construct the local regular solution approximation, which relies on evaluating the chemical potentials and Gibbs curvature matrix at a reference state. Excess solution properties are extracted from this matrix by subtracting off for expected contributions from an ideal solution. The chemical potential and Gibbs curvature contributions of the regular solution are fortunately linear in the excess mixing parameters W_{ij} , enabling linear least-squares inference of their appropriate values. This approach yields a well-behaved perturbation model whose composition-dependence captures the dominant logarithmic and linear terms. When applied to complex multi-site solution phases, this approximation must be carried out in the augmented dependent-species space (rather than simple independent components) to retain meaningful ideal mixing terms.

[NOTE-ASW: Add mathematical details for local regular solution approximation (POST)]

2.4.3 Extrapolation benchmarks for approximate solution models

[NOTE-ASW: Insert discussion and quantitative support for efficacy of solution approximations, especially local regular solution.]

[NOTE-ASW: TODO-> Outline section]

2.5 Final optimization using least squares refinement

[NOTE-ASW: Revisit least squares approach with working local regular solution in mind. Can we linearize that model, pulling approximate ideal mixing into the least squares framework??]

We have thus-far presented an iterative progressive convergence method that ensures relatively fast and accurate results, especially when coupled with compositional solution approximations. But low-level routines like MEXQAL are called thousands of times (or more) in a typical large-scale simulation application, such as tracking an evolving phase assemblage or calibrating a new thermodynamic model requiring many repeat calculations with slightly altered parameter values. With such highly repetitive applications comprising the primary use-cases for this algorithm, speed is of the utmost importance.

We are therefore motivated to develop a modified procedure that linearizes the compositional update procedure to obtain further performance gains. This section focuses on approximating the chemical potential expression for the phase by shifting to logarithmic-composition space, enabling the use of highly optimized linear least-squares methods

to update the composition with each iteration, completely replacing the inner iteration loop discussed previously with a single optimal-update step. While this is in principle fairly straightforward, additional complications arise surrounding the use of degenerate species (rather than independent components)—as required for complex multi-site solution phases like clinopyroxene and spinel—and the necessary mathematical transformations are also derived and presented below.

[NOTE-ASW: the least-squares approach is really only useful in the final stages of the convergence loop, reducing the last dozen iterations or so to only 2 or 3. Include this in section intro.]

[gist: Equilibrium and normalization constraints are linearized by transforming to fraction composition adjustments; enables least-squares refinement of phase] In this section, we show how least squares minimization can provide an approximate coupled solution for composition and affinity of a solution phase in a single step. To apply least squares methods, we must first linearize the exchange equilibrium condition with respect to fractional changes in composition. **[NOTE-ASW: awkward...]** Translating the curvature matrix to log-space results in an expression whose compositional-dependence is linear in log-composition terms ($\log X_i$). Since metastable equilibrium holds simultaneously for each component of the solution, we obtain an approximate system of equations with N constraints:

$$A + \sum_j \frac{\partial \mu_i}{\partial \log X_j} \cdot \delta X_j = \hat{\mu}_i - \mu_i^{\text{ref}} \quad (6)$$

where $\delta X_j = \log X_j - \log X_j^{\text{ref}}$ is the logarithmic change in the composition with respect to the reference, approximately equal to fractional changes as long as they are small. Yet there remains an additional unknown, the saturation affinity of the phase A, which requires a final normalization constraint (or closure condition) for the composition, ensuring that the sum of all components equals 1. Though we cannot perfectly express this in log-space, it can be approximated as:

$$\sum_i X_i^{\text{ref}} \cdot \delta X_i = 0 \quad (7)$$

where the δX_i provide fractional adjustments to the reference composition, and since the reference composition is properly normalized, these changes must all sum to zero to ensure that the final result remains (approximately) normalized.

Combined, the two expressions above provide a system of N+1 linear equations constraining the value of N+1 unknowns (δX_i and A). They can thus be rapidly solved using any standard linear solver.

[NOTE-ASW: summarize this paragraph in intro to least-squares section above. It should be totally clear to reader from the outset that the purpose of the least-squares approach is to speed up the final convergence stage, polishing the result till completion.] It is important to recognize that the inherent assumption that fractional composition adjustments are small will hold true near the end of the convergence process, but may be violated quite strongly at the beginning, depending on the quality of the initial guess. In practice, we find that this poses no problem for most simple solutions (like feldspar or silicate liquid), but does indeed cause serious errors for more complex solutions involving many dependent species, especially when a reasonable initial compositional guess is unknown and we are forced to rely upon the endmember-based initialization method described in Ghiorso (2013). To address this, we test the optimal suggested composition update in each iteration to verify that none of the suggested magnitudes exceed a sensible cutoff (like ~0.01). If any of the solution values exceed this threshold, indicating an unacceptably large composition adjustment that violates the assumptions of the least squares approach, then the suggestion from the direct update method is utilized for that iteration.

3 Results

[NOTE-ASW: Old title #Benchmark accuracy and speed for geologically-relevant phases]

[NOTE-ASW: Add overview of benchmark tests - basic tests, geo]

3.1 Computational efficiency for simple solution phases

[NOTE-ASW: Presents results for simple solution phases like liquid and feldspar(?)] To test this improved algorithm, we can use a variety of thermodynamic solution phases. This will be explored in detail in a

future post. But for now, we can see the immediate impact by testing it on the MELTS liquid phase, which has the advantage of being a rather difficult test requiring simultaneous convergence in 15 different components. In this case, we find that MEXQAL performs extremely well when prototyped using Python code. Whereas the original saturation algorithm of Ghiorso (2013) takes ~280 ms to converge using a “cold-start” (with no initial compositional guess), MEXQAL converges in only 8 ms. These times primarily reflect the number of iterations required as well as the sizable overhead occupied by the objective-c/python bridge software required to evaluate the phase properties. We therefore expect that an implementation of the new algorithm relying directly on cython-wrapped C code will be then again more than an order of magnitude faster.

[gist: emphasize distinction between cold, warm, and hot-start applications for later reference in Results section] To be completely transparent, it is thus useful to consider the number of iterations required for convergence. Here, we focus on the typical use-case, termed a “warm-start,” where a reasonable initial compositional guess is provided and we are simply refining that guess in response to changing conditions. Using MEXQAL without any derivative information, the algorithm converges for a warm-start in ~12-15 iterations. Utilizing the curvature approximation, this decreases to only two iterations of the expensive outer convergence loop.

3.2 Realistic benchmarks for geologically-relevant phases

[gist: Geo-benchmarks for near-solidus equilibrium MORB assemblage] To demonstrate its general power and utility, we test the MEXQAL method for determining saturation affinities and compositions on a variety of geologically-relevant phases. These phases span the range from small simple regular solution models (like ternary feldspar) up to large reciprocal solutions with multiple ordering parameters (like spinel). Typical igneous phases that we selected for testing are presented here—roughly in order of increasing complexity:

- **Feldspar:** regular ternary solution
- **Silicate Liquid:** regular 14-component solution
- **Clinopyroxene:** ordered solution with 7-components & 2 ordering parameters
- **Spinel:** complex ordered solution with 5-components & 4 ordering parameters

For each of these phases, we perform a set of tests utilizing geologically-realistic phase compositions. To sample a region of composition space that is important for geological processes, we simulate near-complete (~95%) batch crystallization of MORB-composition liquid at 1300 K and 1 kbar, relatively close to the solidus. The resulting simulated phase compositions represent realistic compositions for all of the simulated phases, and provide a good indication of expected performance during typical theoretical geochemistry simulations. The simulated compositions used for this study are given in the table below:

Table 1: **Primitive MORB Liquid & Crystallized Mineral Phases**

Oxide	wt%	Endmember	mol%	Endmember	mol%
Liquid		Feldspar		Clinopyroxene	
SiO ₂	48.68	albite	34.98	diopside	45.2
TiO ₂	1.01	anorthite	64.88	clinoenstatite	17.1
Al ₂ O ₃	17.64	sanidine	0.14	hedenbergite	23.7
Fe ₂ O ₃	0.89	—		Al-buffonite	8.0
Cr ₂ O ₃	0.0425	Spinel		buffonite	-1.1
FeO	7.59	chromite	2.3	essenite	4.9
MgO	9.1	hercynite	-30.1	jadeite	2.2
CaO	12.45	magnetite	24.9		
Na ₂ O	2.65	spinel	36.3		
K ₂ O	0.03	ulvospinel	66.6		
P ₂ O ₅	0.08				
H ₂ O	0.2				

- (Note that negative components are required to represent some compositions.)

3.3 Accuracy benchmarks for igneous phases

[NOTE-ASW: full update required based on final benchmarks. Descriptions here are now outdated.]

[gist: Complex phases must be translated to dependent species space] Using the phase compositions given in the table above, we are then prepared to test out the efficacy of the MEXQAL method for this set of solution models of varying complexity. Tests are performed using these compositions by directly calculating the true chemical potentials for each phase, evaluated at this representative equilibrium state. For complex phases with multi-site substitution, these phase compositions and corresponding chemical potentials are converted into the larger degenerate space of species mol fractions. Though this description is redundant, as discussed above, it is critical to MEXQAL approach, enabling the representation of an ideal mixing contribution. The MEXQAL method is then applied to these chemical potentials to invert for a nominally unknown composition. Convergence is assessed by comparing the inferred composition with the true input values, as well as verifying that the estimated phase affinity falls close to zero as expected for phases in the equilibrium assemblage.

[gist: quantify convergence level] Using this approach for all four of the phases described above, we find excellent recovery of the known correct phase composition using the MEXQAL algorithm. It is important to note that translation into dependent species space is critical to the success of these tests, and progress is impossible without that step. For 3 of the 4 phases explored, the algorithm converges on a composition with a corresponding affinity of less than .001 J, fully consistent with being a member of the equilibrium assemblage. Notably, this includes the most complex phase of the set, spinel; even with four ordering parameters, MEXQAL was able to retrieve the correct composition to better than 2 parts in 10^4 for all components. The only phase to show some minor difficulties is clinopyroxene, which converged with an affinity excess of ~ 10 J, and an error in composition at the 1% level for the buffonite endmember. This level of convergence is passable for many applications, but likely results merely from a need to use multiple initial guesses for the MEXQAL method (starting from each dependent species).

3.4 Speed benchmarks for complex phases

4 Discussion

4.1 Comparison with existing algorithms

- PerpleX
- Ghiorso2013 saturation state algo

4.2 Liquid line of decent calculation

- inject this method into standard MELTS calculation
- demonstrates efficacy on standard geo-calculation
- include timing comparison w/ previous Ghiorso (2013) method

4.3 Immiscible tie-line calculation

- must be accompanied by appendix section
- based on simplified deCapitani algorithm

4.4 Mapping miscibility gaps in composition space

- gap finder algorithm
- present results for Feldspar ternary
- present results for Cpx quadrilateral

4.5 exsolve Fe-rich liquids

4.6 Fsp geothermometer

5 Conclusion

Acknowledgments

We thank the following people for useful conversations: Robert Myhill

References

Ghiorso, Mark S. 2013. “A Globally Convergent Saturation State Algorithm Applicable to Thermodynamic Systems with a Stable or Metastable Omni-Component Phase.” *Geochimica Et Cosmochimica Acta* 103 (February): 295–300. <http://www.sciencedirect.com/science/article/pii/S0016703712006667>.

6 Appendix

6.1 Projecting composition into dependent species space

[gist: present conversion to dependent species space including linear and curvature terms.] For complex solutions with multi-site mixing, it is crucial to transform the composition to dependent species space, which raises additional mathematical challenges associated with deriving properties of the dependent species from those of the endmember components. **[NOTE-ASW: Perhaps use \sim mark for endmembers rather than species, enabling rest of text to be correct when applied to species (for simple solutions, endmember components are the only species).]** This transformation rests directly on the species stoichiometry matrix:

$$\tilde{\nu}_{ij} = \partial n_j / \partial \tilde{n}_i$$

which describes how the number of moles of the j^{th} endmember component (n_j) changes in response to changes in the amount of the i^{th} species (\tilde{n}_i). The chemical potentials of the species are then given simply by the stoichiometrically-weighted average of the endmember chemical potentials:

$$\tilde{\mu}_i = \sum_j \tilde{\nu}_{ij} \mu_j$$

Likewise, we must also find a transformation for the Gibbs curvature matrix, which describes the linear dependence of chemical potentials on changes in molar composition, but it does so in endmember component space. When solving problems in the larger degenerate compositional space of dependent species, we must mathematically inflate this matrix to describe how the chemical potentials of these species change as their quantity is altered.

$$\frac{\partial \tilde{\mu}_i}{\partial \tilde{n}_l} = \sum_k \sum_j \tilde{\nu}_{ij} \frac{\partial \tilde{\mu}_j}{\partial \tilde{n}_k} \tilde{\nu}_{kl}^T$$
$$\frac{\partial \tilde{\mu}}{\partial \tilde{n}} = \tilde{\nu} \cdot \frac{\partial \tilde{\mu}}{\partial \tilde{n}} \cdot \tilde{\nu}^T$$

where ν^T is the transpose of the stoichiometry matrix (oriented with species as columns rather than rows). **[NOTE-ASW: Majorly unclear. Perhaps equation needs to show intermediate steps? $\mathbf{d}_{\mu}/\mathbf{d}\mathbf{n} = \text{sum}(\mathbf{d}\sim\mu/\mathbf{d}\mathbf{n} \mathbf{d}\mathbf{n}/\mathbf{d}\sim\mathbf{n})$** This transformation is derived by applying the stoichiometry matrix once to calculate the linear dependence of species on the molar endmember composition, and then a second time to calculate their dependence on the molar abundance of the species themselves. Since the stoichiometry matrix is unchanging and fixed by the composition of each species, this new Gibbs curvature matrix for dependent species retains all the linearity of the original in endmember component space, thereby allowing an identical least squares approach to solve for dependent species abundances.

6.2 Basic composition-perturbation models for solution phases

The simplest local approximations for chemical solutions are the physically-motivated ideal solution and the purely mathematical linear solution models. Because the ideal solution is based on the assumption of ideal mixing, it completely ignores any chemical interaction energies and thus generally performs rather poorly for most real-world solutions. In stark contrast, the linear solution model is a purely mathematical representation that has no physical basis, but shows much improved performance as it empirically accounts for (linear) compositional variation in the chemical potentials. Together, these two rather basic yet important models provide the foundation for a much more accurate approximation, the local regular solution, which we will later explore in detail.

The ideal solution model represents a thermodynamic baseline (or null hypothesis) for mixed phases; though it is rarely used directly, it forms the foundation of nearly all more complicated models. The fundamental assumption underlying the ideal solution is that the components of the solution behave as energetically independent entities. The resulting mixture thus has an energy given by the weighted average of the components, modified only by the additional entropy of mixing, yielding the simple molar Gibbs energy expression:

$$\bar{G}^{ideal} = \sum_i X_i \mu_i^0 + RTm \sum_i X_i \log X_i \quad (8)$$

where μ_i^0 are the chemical potentials for each pure endmember component, X_i is the mol-fraction of each component, RT is the thermal energy scale, and m is the site multiplicity. The chemical potentials are given by the compositional derivative:

$$\mu_k = \mu_k^0 + RTm \ln X_k$$

which emphasizes how the chemical potential of the pure endmember component is only modified by ideal entropy mixing. Based on this form, the ideal solution can be transformed into a local perturbation model:

$$\Delta\mu_k^{ideal} = \mu_k - \mu_k^{ref} \approx RTm \cdot \ln (X_k/X_k^{ref})$$

where changes in the chemical potentials are attributed entirely to the expected behavior for the entropy of mixing terms.

Conceptually much simpler is the linear solution model, which is simply a first-order Taylor expansion of the chemical potentials with respect to the molar composition of the phase. Unlike the ideal solution, the linear solution thus requires that both the chemical potentials and their compositional derivatives (the Gibbs curvature matrix) be calculated at the reference composition:

$$\Delta\mu_k^{linear} = \mu_k - \mu_k^{ref} \approx \sum_l \frac{d\mu_k}{dn_l} \cdot (X_l - X_l^{ref})$$

where changes in the chemical potentials are assumed to have linear dependence, regardless of the magnitude of the compositional extrapolation. For small perturbations, this assumption is perfectly reasonable, and the linear solution model thus greatly outperforms the ideal solution in this regime. But the lack of logarithmic mixing terms comes to dominate for larger extrapolations, urging us to seek a more accurate model that can capture both empirically observed chemical potential variations and the appropriate form for entropic mixing.

6.3 Key Points (working collection)

[NOTE-ASW: Needs to be paired down/simplified before moving to document frontmatter.]

- The saturation state algorithm of Ghiorso (2013) is extended to enable calculation of metastable exchange equilibria, even in the absence of an omni-component phase, useful for subsolidus and miscibility gap calculations.
- MEXQAL algorithm determines chemical exchange equilibrium for a solution phase; from an externally imposed chemical potential, we can determine phase composition X and saturation affinity A consistent w/ metastable equilibrium.
- Affinity and composition are directly determined using a new iterative update scheme, which is more numerically accurate and faster than the original repeated fractions method of Ghiorso (2013)

- Convergence speed is further improved using derivative information, which enables approximate convergence within an inner loop, and dramatically reducing the number of direct function evaluations, particularly important for complex phases involving ordering.
- Exchange equilibrium conditions can be found approximately by transforming to log-composition space; this yields a set of coupled linear equations, with $\Delta \log X_i$ and A as unknowns, that can be solved using least squares.
- Fractional composition adjustments must be rather small (<0.1), otherwise the accuracy of the approximate normalized composition equation suffers.
- For complex phases, the Gibbs curvature matrix must be remapped to dependent species space, inflating the matrix to describe how species' chemical potentials change with species abundances; this operation relies only on the known stoichiometry of each species.
- Testing of igneous phases under realistic magmatic conditions reveals excellent performance for simple phases like liquid and feldspar, but only modest gains for complex phases like spinel and clinopyroxene.
- Added computational cost of the least squares approach could be avoided for most early iterations by accepting direct compositional update if any components exceed a threshold value.
- Typical igneous phases (liquid, Clinopyroxene, Feldspar, and Spinel) are used to test the MEXQAL method
- Realistic compositions are determined using batch crystallization of primitive MORB liquid
- Complex solutions with multi-site mixing require negative components to reach all of composition space; MEXQAL thus needs to work with dependent species (rather than independent components) to ensure all positive compositional variables, enabling usage of ideal mixing correction ($RTm * \log X_i$)
- Tests for all phases show excellent convergence to the correct compositions using the MEXQAL method
- Only clinopyroxene shows some deviation (at the 10 J affinity level) that needs further exploration
- Ideal solution model assumes non-interacting components in the mixture, leading to a weighted average Gibbs energy modified only by the entropy of mixing.
- Linear solution model provides a basic linear description for solution phases which has greater accuracy for small compositional perturbations, but lacks any awareness of entropic mixing.
- The best aspects of each of these models can be combined to form the local regular solution approximation, which builds a complete regular solution based on phase properties evaluated at the reference composition, allowing greater extrapolation in composition space.