

Winding journey down (in Gibbs energy)

Jeffrey M. Dick

January 1, 2013

1 Introduction

`wjd()` appeared in CHNOSZ version 0.9-9. It implements the steepest descent algorithm for Gibbs energy minimization described by White, Johnson and Dantzig, 1958. This vignette shows some examples of its usage. The help page for `wjd()` in the CHNOSZ documentation has details of the implementation that are not covered in this vignette. The functions demonstrated here are not intended for aqueous solutions or heterogeneous phase equilibria, which are covered in some reviews (e.g. Nordstrom et al., 1979).

The arguments of `wjd()` have default values to simulate the example problem given by White et al., 1958. In their words, “The example taken is the determination of the composition of the gases arising from the combustion of a stoichiometric mixture of hydrazine and oxygen at 3500 °K and a pressure of 750 psi ($\ln P_{\text{atm}} = 3.932$).” There are ten species included in this example. The value of temperature does not appear explicitly in the algorithm, or in the arguments to `wjd()`. Instead, the standard Gibbs energies of the species, at the given temperature, are provided in dimensionless form, i.e. $\Delta G^\circ / RT$. Note that White et al., 1958 use the terminology “free energy” and the notation F° . The term “Gibbs energy” and corresponding notation are used here.

2 Is it winding?

The “winding” in the title refers to the observation that the abundances of species during a Gibbs energy minimization often do not all change in a monotonic fashion. This is because at each iteration a new direction of steepest descent is calculated; this direction is a vector of changes of mole numbers of all species in the system and is subject to mass balance constraints as well as the steepest-descent criterion.

Let’s run `wjd()` with its default settings and save the output.

```
> w <- wjd()
```

What are the most abundant species, and how many iterations were taken?

```
> # the order of species abundance
> oX <- order(w$X, decreasing=TRUE)
> # the stoichiometries of the two most abundant species
> w$A[head(oX,2),]
```

```
      H N O
[1,] 2 0 1
[2,] 0 2 0
```

```
> # the number of iterations
> niter <- length(w$lambda)
> niter
```

```
[1] 6
```


We find that H₂O and N₂ are the most abundant species, after 6 iterations.

Let's track their mole fractions through the iterations. Write a function that returns the mole fractions of these two species, after a specified number of iterations.

```
> iterfun <- function(imax) {  
+   w <- wjd(imax=imax)  
+   x1 <- w$X[oX[1]]/sum(w$X)  
+   x2 <- w$X[oX[2]]/sum(w$X)  
+   return(list(x1=x1, x2=x2))  
+ }
```

Then apply the function over the different numbers of iterations, from 0 (initial conditions) to 6, and plot the values.

```
> sa <- sapply(0:niter, iterfun)  
> par(mfrow=c(1, 2))  
> plot(0:niter, sa[1, ], xlab="iteration", ylab=paste("x", f1))  
> plot(0:niter, sa[2, ], xlab="iteration", ylab=paste("x", f2))
```



A bit of winding: the mole fractions of H₂O and N₂ increase up to the third iteration, but at the fourth (and less so for the fifth and sixth), their mole fractions decrease. This behavior is consistent with a decrease in Gibbs energy at *every* iteration; that can be verified by inspecting the values in the 6-iteration result:

```
> diff(w$F.Y)  
[1] -9.001450e-01 -1.690993e-01 -2.702223e-02 -1.348404e-03 -9.581509e-05  
[6] -3.417140e-06
```

The decrease in Gibbs energy becomes smaller with every iteration, as an equilibrium distribution of species is approached.

3 Is it near the bottom? Testing for equilibrium

3.1 Change in Gibbs energy of the system becomes small

Equilibrium is synonymous with a global Gibbs energy minimum. Becoming convinced that the output from `wjd()` represents a near-equilibrium condition can be difficult. One type of observation that can be

helpful is the amount of change during the iterations of the algorithm. As equilibrium is approached, it makes sense that the fractional decreases in Gibbs energy would become smaller and smaller. This is used as a stopping condition in the current implementation – if the fractional change, relative to the total energy of the system, reaches the value of `Gfrac` given in the arguments the iterations stop. Here are the fractional changes with each iteration:

```
> diff(w$F.Y)/w$F.Y[1:6]
[1] 1.929006e-02 3.555210e-03 5.661133e-04 2.823296e-05 2.006125e-06
[6] 7.154610e-08
```


Only the last value is below the default value of `Gfrac` in `wjd()`.

3.2 Chemical potentials of the elements (from different species representing them) become more uniform

Equilibrium is also synonymous with uniformity of chemical potentials (μ) of the thermodynamic components throughout the system. For the default system in `wjd()`, and perhaps other systems of interest, the elements themselves are a conceivable set of the components. Often, the number and compositions of the species are such that multiple combinations of species satisfy the stoichiometric conditions necessary to be a basis set, and therefore can be used to compute the chemical potentials of the elements. At equilibrium, these different combinations of species would all yield the same chemical potentials of the elements.


There is a supporting function, `element.potentials()`, that computes the chemical potentials of the elements using different eligible combinations of species. It has an option to plot the results. Let's first look at a plot showing the results after 3 iterations.

```
> w3 <- wjd(imax=3)
> ep3 <- element.potentials(w3, plot.it=TRUE)
```



Here's the plot for the default settings of `wjd()` which takes 6 iterations:

```
> ep <- element.potentials(w, plot.it=TRUE)
```



That one shows considerably less deviation than the first plot. The differences don't become zero, but perhaps they are small enough to accept as an operational solution. There is a function, `is.near.equil()`, that provides a logical type result based on the range of chemical potentials of the elements calculated from the different eligible species combinations. It returns TRUE if the maximum difference for any element is below $\mu/RT = 0.01$, but this limit can be changed through the `tol` argument to the function. With the default setting, `is.near.equil()` does pass only 6-iteration result, but by reducing the tolerance, the tests can be made to fail:

```
> is.near.equil(w3)
```

```
is.near.equil: solution has variation of 3.49055498794722 in mu/RT of 0
[1] FALSE
```

```
> is.near.equil(w)
```

```
[1] TRUE
```

```
> is.near.equil(w, tol=0.0001)
```

```
is.near.equil: solution has variation of 0.00726239786434846 in mu/RT of 0
[1] FALSE
```

3.3 Different initial solutions yield similar results

Different guesses for an underdetermined system Although it seems to be the case in the examples shown so far, increasing the number of iterations does not necessarily bring one closer to true equilibrium, since being trapped in a local energy minimum is a possibility. Sometimes local minima can be shown to exist by starting with different mole numbers of species, that still yield the same bulk composition (mole numbers of elements). `guess()` exists to facilitate this sort of investigation. Its input the stoichiometric matrix of species in the system (the default refers to the same system as in `wjd()`) and the bulk composition; it outputs one set or a series of sets of mole numbers of species that satisfy the bulk composition constraints. In general, this is an underdetermined problem, so there are usually an infinite number of possible solutions.

Using the default method (“stoich”) a single guess or multiple guesses are returned based on stoichiometric constraints (see the help page for details as well as the comments in the function code itself). If the “central” method is chosen instead, the single guess returned is the central solution from linear optimization using the package `limSolve`. Let’s compare the guesses made by `guess()` with the default (from [White et al., 1958](#)) initial composition in `wjd()`.

```
> as.list(args(wjd))$Y
c(0.1, 0.35, 0.5, 0.1, 0.35, 0.1, 0.1, 0.1, 0.1, 0.1)
> Y1 <- guess(method="stoich")
> Y1
[1] 0.012 0.002 0.990 0.992 0.002 0.002 0.002 0.002 0.002 0.002
> Y2 <- guess(method="central")
> Y2
[1] 0.10065 0.15110 0.48210 0.05075 0.07555 0.53255 0.26560 0.05070 0.05060
[10] 0.10040
```

Now let’s run the minimization using the different guesses, count the numbers of iterations, and test them for equilibrium.

```
> wY1 <- wjd(Y=Y1)
> niterY1 <- length(wY1$lambda)
> niterY1
[1] 9
> is.near.equil(wY1, tol=0.0001)
is.near.equil: solution has variation of 0.000641091319222298 in mu/RT of 0
[1] FALSE
```


Using the first guess generated using the “stoich” method there were 3 more iterations than using the default initial solution in `wjd()`.

```
> wY2 <- wjd(Y=Y2)
> niterY2 <- length(wY2$lambda)
> niterY2
[1] 8
> is.near.equil(wY2, tol=0.0001)
[1] TRUE
```

The initial guess generated using the “central” method lead to 2 more iterations than encountered using the argument defaults in `wjd()`, and the algorithm converged to an even lower variability in chemical potentials of the elements than using the “stoich” method. The differences in convergence could be coincidental, and the “stoich” method might be preferable for general usage because it allows multiple guesses to be tested automatically (see below).

Do the different initial guesses actually give similar results?

```
> plot(1:10, w$X, xlab="species", ylab="mole fraction")
> points(1:10, wY1$X, pch=0)
> points(1:10, wY2$X, pch=2)
```



At that scale it's fairly difficult to tell them apart.

Comparing many different guesses Let's test for closeness to equilibrium of all of the guesses available from `guess()` for the default set of species and bulk composition. First generate all of the guesses. Note that `guess()` returns NA for compositionally eligible combinations of species that have some negative mole numbers, so we filter those out.

```
> Ys <- guess(iguess=NULL, method="stoich")
> # total number of species combinations
> length(Ys)

[1] 76

> # species combinations that have all-positive mole numbers
> iYs <- which(!is.na(Ys))
> nguess <- length(iYs)
> nguess

[1] 39
```

Now let's run the minimization using each of those guesses, and test if each one `is.near.equil()`:

```
> sapply(iYs,function(i) is.near.equil(wjd(Y=Ys[[i]])))

[1] TRUE TRUE
[16] TRUE TRUE
[31] TRUE TRUE TRUE TRUE TRUE TRUE TRUE TRUE TRUE
```

Looks like they are. What if we're really picky though and want to make sure the chemical potentials of the elements are very well-defined? We can decrease the tolerance of `is.near.equil()`:

```
> sapply(iYs,function(i) is.near.equil(wjd(Y=Ys[[i]]), tol=0.0001))
```

```

is.near.equil: solution has variation of 0.000641091319222298 in mu/RT of 0
is.near.equil: solution has variation of 0.000383472339134272 in mu/RT of 0
is.near.equil: solution has variation of 0.00034678256681886 in mu/RT of 0
is.near.equil: solution has variation of 0.000231788387019094 in mu/RT of 0
is.near.equil: solution has variation of 0.00827251816836849 in mu/RT of 0
is.near.equil: solution has variation of 0.000651868600300531 in mu/RT of 0
is.near.equil: solution has variation of 0.000383472339121838 in mu/RT of 0
is.near.equil: solution has variation of 0.000997071906226665 in mu/RT of 0
is.near.equil: solution has variation of 0.000921014280095278 in mu/RT of 0
is.near.equil: solution has variation of 0.000481312160735214 in mu/RT of 0
is.near.equil: solution has variation of 0.000316491422417187 in mu/RT of 0
is.near.equil: solution has variation of 0.000216269345349573 in mu/RT of 0
is.near.equil: solution has variation of 0.00231700216697739 in mu/RT of 0
 [1] FALSE TRUE FALSE TRUE TRUE TRUE TRUE TRUE FALSE FALSE TRUE TRUE
[13] FALSE TRUE TRUE TRUE FALSE TRUE FALSE TRUE TRUE FALSE FALSE TRUE
[25] TRUE TRUE TRUE TRUE TRUE TRUE TRUE TRUE TRUE FALSE TRUE TRUE FALSE
[37] TRUE FALSE FALSE

```

Well some of the tests failed. But maybe this is because we have stopped iterating too soon. Let's iterate until a we have smaller change in Gibbs energy of the system:

```

> sapply(iYs, function(i) {
+   is.near.equil(wjd(Y=Ys[[i]], Gfrac=1e-9), tol=0.0001)
+ })

```

```

is.near.equil: solution has variation of 0.000316491422417187 in mu/RT of 0
 [1] TRUE TRUE
[13] TRUE TRUE
[25] TRUE TRUE
[37] TRUE TRUE TRUE

```

We're back to passing all but one of the equilibrium tests based on uniformity of chemical potentials of the elements. What are the standard deviations of the resulting species abundances?

```

> Xs <- sapply(iYs, function(i) wjd(Y=Ys[[i]], Gfrac=1e-9)$X)
> apply(Xs, 1, sd)

```

```

 [1] 2.913889e-10 1.441534e-09 2.679858e-09 2.386238e-08 1.188050e-08
 [6] 4.017052e-09 3.779520e-10 1.496856e-10 8.115481e-10 3.224370e-09

```

Based on this exercise, starting from different initial species abundances for the same bulk composition, and arriving at a similar near-equilibrium solution, it seems likely that we are near the global minimum.

(In fact, for ideal reactions in closed chemical systems there is only a unique solution (Othmer, 1976). Therefore, the algorithm should converge on the same result for any initial vector of species mole numbers with the same bulk composition.)

4 An example from the literature: prebiological atmospheres

We will try to reproduce a subset of the calculations of equilibrium between carbon-containing compounds in model prebiological atmospheres reported by Dayhoff et al. (1964), using standard Gibbs energies from Dayhoff et al. (1967). The standard Gibbs energies shown in Figure 2 of Dayhoff et al. (1967) for the selected compounds listed in Table 1 of Dayhoff et al. (1964) are provided with CHNOSZ. First let's have a look at what species there are:

```

> # read formulas and Gibbs energies
> file <- system.file("extdata/thermo/DLEN67.csv", package="CHNOSZ")
> dlen <- read.csv(file, stringsAsFactors=FALSE, row.names=1)
> t(dlen[, 1, drop=FALSE])

      water carbon dioxide nitrogen methane ammonia carbon monoxide ethane
formula "H2O" "CO2"           "N2"      "CH4"  "NH3"   "CO"           "C2H6"
      formic acid propane hydrogen cyanide methylamine butane glycine
formula "CH2O2"  "C3H8"  "HCN"           "CH5N"   "C4H10" "C2H5NO2"
      benzene hexane naphthalene anthracene octane
formula "C6H6"  "C6H14" "C10H8"      "C14H10" "C8H18"

```

Now let's write some code to define the system. The bulk composition in C:H:O space has 40 percent oxygen, a percentage of carbon given by x_C , and hydrogen as the remainder; the C:N ratio is taken to be 1:1 (Dayhoff et al., 1964). The temperature is 500 K and the pressure is 1 atmosphere.

```

> # turn formulas into a stoichiometric matrix
> A <- i2A(dlen$formula)
> # assemble Gibbs energies/RT at 500 K
> G0.RT <- 1000 * dlen$G500 / thermo$opt$R / thermo$opt$Tr
> # a function to minimize Gibbs energy for system with
> # given mole fraction of carbon (xC)
> min.atmos <- function(xC) {
+   # the bulk composition C:H:N:O
+   B <- c(xC, 100-40-xC, xC, 40)
+   # guess the initial composition
+   Y <- guess(A, B)
+   w <- wjd(A=A, G0.RT=G0.RT, Y=Y, P=1, imax=90, Gfrac=1e-14)
+   if(!is.near.equil(w)) cat(paste("not near equilibrium for xC=", xC, "\n"))
+   return(w)
+ }

```

Notice the increase in $imax$ and the decrease in $Gfrac$ from the defaults! These changes are needed to get closer to equilibrium (and we'll know we're not there if a message is shown). What does the prebiological atmosphere look like with 15% (not counting nitrogen) carbon?

```

> sort(min.atmos(15)$X, decreasing=TRUE)

      H2O      CO2      N2      CH4      NH3      CO
1.624999e+01 1.187500e+01 7.499996e+00 3.124996e+00 8.375716e-06 1.712505e-06
      C2H6      CH2O2      C3H8      CH5N      HCN      C4H10
5.171341e-10 1.768497e-14 2.299958e-18 2.507394e-20 1.164070e-20 3.627367e-26
      C2H5NO2      C6H6      C6H14      C8H18      C10H8      C14H10
2.098217e-33 5.675008e-36 1.317368e-42 5.118505e-59 2.676202e-59 4.102681e-84

```

Well that's exciting! The order of abundances of H_2O , CO_2 , CH_4 , NH_3 , CO , ethane (C_2H_6), formic acid (CH_2O_2) and so on matches the order on the left side of Figure 2 in Dayhoff et al. (1964). How about changing the carbon content?

```

> xCs <- seq(8, 47, 3)
> Xs <- sapply(xCs, function(xC) min.atmos(xC)$X)

```

```


is.near.equil: solution has variation of 3.07334312115114 in mu/RT of 0
not near equilibrium for xC= 14

```

```

> # normalize the mole numbers to mole fractions
> Xs <- t(t(Xs)/colSums(Xs))
> plot(-10, 0, xlim=c(0, 55), ylim=c(-25, 1), xlab="mole percent C", ylab="log10 mole fraction")
> for(i in 1:nrow(Xs)) lines(xCs, log10(Xs[i, ]))
> text(48, log10(Xs[, length(xCs)]), dlen$formula, adj=0)
> text(7, log10(Xs[, 1]), dlen$formula, adj=1)

```



We triggered a single “not near equilibrium” message, but overall it seems well behaved. And, as expected, it is similar to Fig. 2 of [Dayhoff et al. \(1964\)](#), with a major crossing of curves at about 28% carbon, together with an increase in aromatic compounds (e.g. naphthalene, anthracene) going toward higher carbon content. Unlike the figure in [Dayhoff et al. \(1964\)](#), there appears to be a second major crossing of curves at about 43% carbon, corresponding to a rise in CO. (For a higher resolution, try setting `xCs <- seq(8, 47, 1)`.)

5 Running Down: Using a Thermodynamic Database

So far the examples shown here have been based on the chemical system defined by the default settings for the arguments to the functions. What if we’re interested in a different system? It can be rather tedious to manually construct the stoichiometric matrices and vectors of standard Gibbs energies of the species. The function `run.wjd()` pulls the compositional and thermodynamic data for requested species from the thermodynamic database and then calls `wjd()`.

5.1 At a single temperature

Let’s start by finding the indices in the thermodynamic database of some liquid alkanes and aromatic compounds:

```

> alkanes <- c("n-hexane", "n-heptane", "n-octane", "n-nonane")
> ialk <- info(alkanes, "liq")
> aromatics <- c("naphthalene", "anthracene", "phenanthrene", "pyrene")
> iaro <- info(aromatics, "liq")

```

Let's find the equilibrium distribution of species for a bulk composition corresponding to a single mole of each species.

```
> waa <- run.wjd(c(ialk, iaro), Y=rep(1, 8))
```

```
subcrt: 8 species at 298.15 K and 1 bar
```

```
> waa$elements
```

```
  [,1] [,2] [,3] [,4] [,5] [,6] [,7] [,8] [,9] [,10] [,11]
C   84   84   84   84   84   84   84   84   84   84   84
H  106  106  106  106  106  106  106  106  106  106  106
```

```
> is.near.equil(waa)
```

```
is.near.equil: solution has variation of 0.301276003581621 in mu/RT of H
[1] FALSE
```

The `waa$elements` shows that the bulk chemical composition in the Gibbs energy minimization was $C_{84}H_{106}$ ($H/C = 1.261905$). That `FALSE` means the chemical potentials of the elements calculated from different combinations of species differ by more than 0.01, the default setting of `tol` in `is.near.equil()`. We can make it `TRUE` by running more iterations:

```
> waa <- run.wjd(c(ialk, iaro), Y=rep(1, 8), imax=12, Gfrac=1e-14)
```

```
subcrt: 8 species at 298.15 K and 1 bar
```


```
> is.near.equil(waa)
```

```
[1] TRUE
```

Let's plot the abundances of the species:

```
> bp <- barplot(waa$X, ylab="moles")
```

```
> text(bp, rep(0.2,8), c(alkanes, aromatics), srt=90, adj=0)
```



Now let's calculate the chemical potentials of the elements (dimensionless values, μ/RT) ...

```
> ep <- equil.potentials(waa)
> print(ep)
```

```
      C      H
5863.434 -2610.552
```

... and the corresponding logarithms of chemical activities of a set of basis species:

```
> basis(c("graphite", "H2"), c("cr", "gas"))
```

```
  C H ispecies logact state
C  1 0    1937      0   cr
H2 0 2    3090      0   gas
```

```
> basis.logact(ep)
```

```
      C      H2
4.297936 -3.827104
```

5.2 Bulk composition instead of moles of species

run.wjd() can be given B (bulk chemical composition), instead of Y (initial numbers of moles of species). When this happens, the function tests the different initial compositions generated by guess() until one is found that is within the specified tolerance of chemical potentials of elements. Here's an example:


```
> waa <- run.wjd(c(ialk, iaro), B="C100H70")
```

```
subcrt: 8 species at 298.15 K and 1 bar
```

```
run.wjd: got within tolerance on initial solution 1 of 7
```

```
> bp <- barplot(waa$X, ylab="moles")
```

```
> text(bp, rep(0.2,8), c(alkanes, aromatics), srt=90, adj=0)
```



Compared with the previous example, here the lower $H/C = 0.7$ defines a more oxidized system, so the increase in aromatic content is expected. As before, we can calculate the chemical potentials of the elements and of a corresponding set of basis species:

```
> print(ep <- equil.potentials(waa))
```

```
      C      H
5964.177 -2723.926
```

```
> basis.logact(ep)
```

```
      C      H2
4.371781 -3.993311
```

6 Document Information

Revision history

- 2012-01-01 Initial version
- 2012-06-16 Running Down using alkanes and aromatics
- 2012-09-20 Add example from Dayhoff et al.

R session information

```
> sessionInfo()
```

```
R version 2.15.2 (2012-10-26)
```

```
Platform: x86_64-slackware-linux-gnu (64-bit)
```

```
locale:
```

```
[1] LC_CTYPE=en_US      LC_NUMERIC=C          LC_TIME=en_US
[4] LC_COLLATE=C        LC_MONETARY=en_US    LC_MESSAGES=en_US
[7] LC_PAPER=C          LC_NAME=C             LC_ADDRESS=C
[10] LC_TELEPHONE=C      LC_MEASUREMENT=en_US LC_IDENTIFICATION=C
```

```
attached base packages:
```

```
[1] stats      graphics  grDevices  utils      datasets  base
```

```
other attached packages:
```

```
[1] limSolve_1.5.3 MASS_7.3-22  lpSolve_5.6.6  quadprog_1.5-4 CHNOSZ_0.9-9
```

```
loaded via a namespace (and not attached):
```

```
[1] tools_2.15.2
```

References

- M. O. Dayhoff, E. R. Lippincott, and R. V. Eck. Thermodynamic Equilibria In Prebiological Atmospheres. *Science*, 146(3650):1461–1464, December 1964. doi: [10.1126/science.146.3650.1461](https://doi.org/10.1126/science.146.3650.1461).
- M. O. Dayhoff, E. R. Lippincott, R. V. Eck, and G. Nagarajan. Thermodynamic Equilibrium in Prebiological Atmospheres of C, H, O, N, P, S, and Cl. Report SP-3040, National Aeronautics and Space Administration, Washington, D. C., 1967. URL <http://ntrs.nasa.gov/search.jsp?R=19670017966>.

- D. K. Nordstrom, L. N. Plummer, T. M. L. Wigley, T. J. Wolery, J. W. Ball, E. A. Jenne, R. L. Bassett, D. A. Crerar, T. M. Florence, B. Fritz, M. Hoffman, Jr. Holdren, G. R., G. M. Lafon, S. V. Mattigod, R. E. McDuff, F. Morel, M. M. Reddy, G. Sposito, and J. Thraikill. A comparison of computerized chemical models for equilibrium calculations in aqueous systems. In E. A. Jenne, editor, *Chemical Modeling in Aqueous Systems: Speciation, Sorption, Solubility, and Kinetics*, Series 93, pages 857 – 892. American Chemical Society, 1979. URL <http://www.worldcat.org/oclc/4591313>.
- H. G. Othmer. Nonuniqueness of Equilibria In Closed Reacting Systems. *Chem. Eng. Sci.*, 31(11):993–1003, 1976. doi: [10.1016/0009-2509\(76\)87020-0](https://doi.org/10.1016/0009-2509(76)87020-0).
- W. B. White, S. M. Johnson, and G. B. Dantzig. Chemical equilibrium in complex mixtures. *J. Chem. Phys.*, 28(5):751 – 755, 1958. ISSN 0021-9606. doi: [10.1063/1.1744264](https://doi.org/10.1063/1.1744264).