

Isothermal titration calorimetry: feeling the heat

Gemma Harris and David J. Scott

(Research Complex at Harwell and University of Nottingham, UK)

Isothermal titration calorimetry (ITC) has become one of the mainstays of determining the affinity of small molecules with macromolecules. As such there has been a resurgence in interest in the underlying processes that contribute to a 'good' ITC experiment. In this article we discuss the thermodynamic origins of the data seen in an ITC experiment, some of the background theory and exemplars of what good data look like. We also discuss the origins of more complex phenomena seen with ITC, and how these may be interpreted.

Introduction

If you really want to torture first-year undergraduates in the biosciences, then there is no better place to start than a good solid introduction to thermodynamics. This usually begins a frank conversation about pistons and steam engines, moving through to something to do with engines called Carnot, before finally introducing entropy which appears to be simultaneously 'disorder', 'not really disorder, but something like it' or apparently a different flavour of heat from enthalpy (which has something to do with chemical bonds, perhaps). An exam of some sort subsequently ensues: relief then breaks out when most will get a mark just over the necessary threshold to continue, and the subject—much like an embarrassing older relative—is never spoken of again in polite company.

The substance of the above paragraph is clearly a parody, but the tone will likely resonate for many of those tasked with teaching basic chemistry to biosciences students. Thermodynamics is one of the crowning achievements of 19th century science, yet we are expected to simplify discussion of the theme down to a few handy catchphrases and analogies that can be subject to testing by examination.

Despite this scepticism, regarding the outset of bioscience education, the increasing use of isothermal titration calorimetry (ITC) represents a key technique in characterizing biomolecular associations and should therefore be a cause of some celebration. Indeed, ITC is now a standard technique in many laboratories across the world for determining key thermodynamic parameters associated with ligand-binding events. This has brought thermodynamics back front-and-centre in many studies of biomacromolecules. The basis of the technique is to accurately measure the heat produced or taken up upon the binding of one molecule to another: typically, this has been the binding of a small molecule to a protein, and as such is used to determine affinity in a rigorous fashion. Added to this, quantities such as enthalpy, entropy and free energy can also be determined along with the stoichiometry.

In this feature, we are going to concentrate on what a 'good' ITC experiment looks like, how you generate a 'typical' ITC thermogram and what can go wrong with the experiment. We are going to assume that the ligand is a small molecule with good solubility in an aqueous buffer, and the macromolecule in the cell is a protein with a single binding site for the ligand. This will give rise to what could be referred to as a 'standard ITC thermogram' as detailed in Figure 2B, and as seen in numerous publications over the years. There are many very good reviews on what you can do with the derived thermodynamic values and how they can be interpreted structurally, which are beyond the scope of this article, and the interested reader is referred to them.

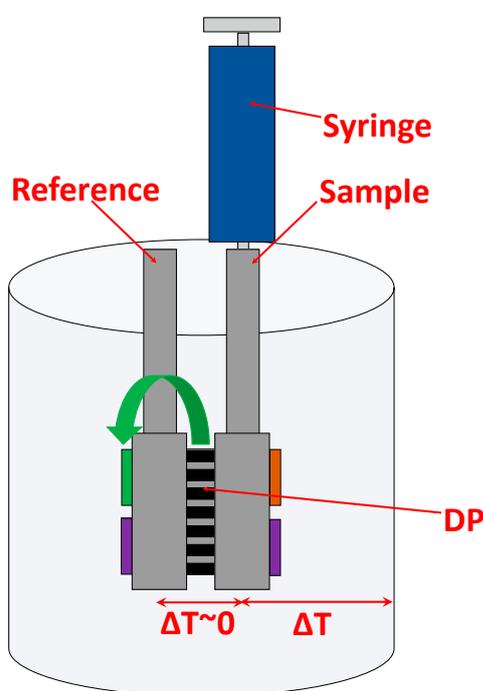


Figure 1. Schematic of an ITC instrument. The differential power between the sample and reference cell is measured and adjusted throughout the experiment to maintain $\Delta T = \sim 0$.

█ Reference heater
█ Sample heater
█ Main heater
 DP - Differential power
 ΔT - Temperature difference

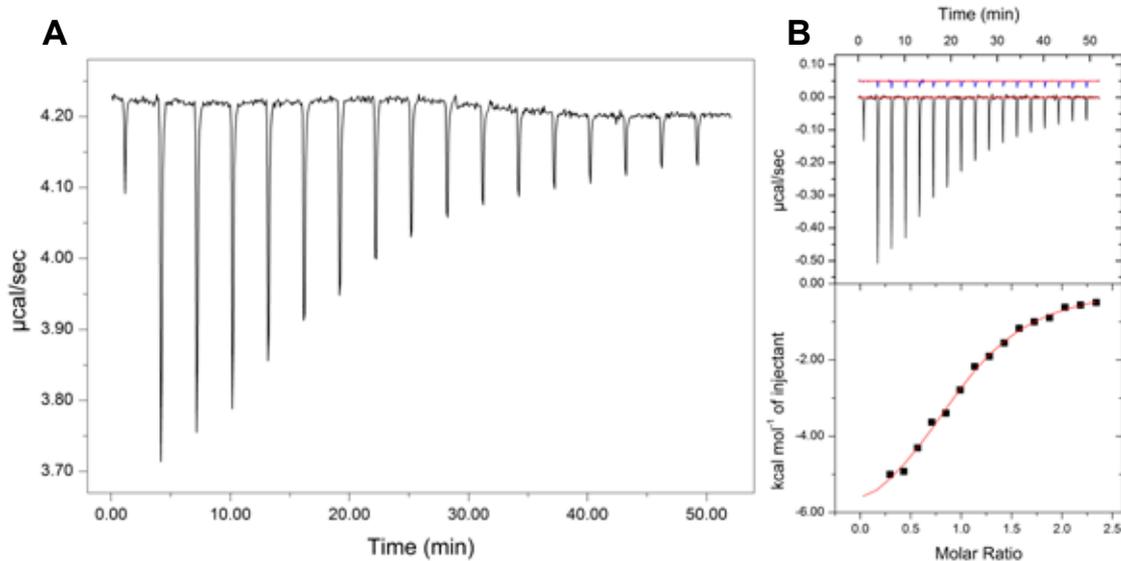


Figure 2. Examples of experimental ITC data. **A)** Raw ITC data. **B) Top panel:** baseline corrected thermograms. The experimental and control data are shown in black and blue, respectively. The baselines are indicated in red. **Bottom panel:** the integrated heats upon injection (black squares) and the data fit (red line) after subtraction of the control data.

What happens during an ITC experiment?

The basic principle of all ITC experiments is the titration of one component into a solution of another (see Figure 1). The heat generated is then quantified by measuring the amount of feedback current required to keep the ITC cell the same temperature as a reference cell. This differential method is key to the technique as this current can be measured with very high accuracy and hence very small temperature differences can be defined with good precision.

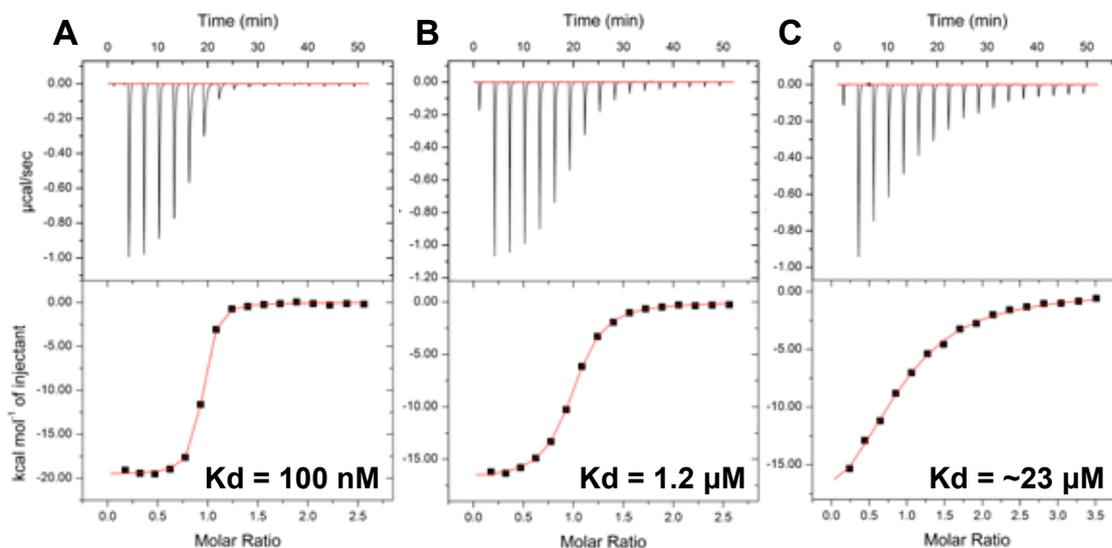
Practically, this means that a 200 μl solution containing approximately 10 μM of protein (approximately 0.3–1 mg/ml) can give a decent ITC trace upon titration with a ligand at 100 μM in the ITC syringe as long as there is sufficient heat generated upon binding (usually >2.5 kJ.mol). These are quantities of protein that are within the range of many labs to produce by standard bacterial over-expression, and the protein can be recovered at the end of the experiment and used again, making ITC an attractive option.

Figure 2A shows the result of an experiment titrating a small molecule ligand into a solution of protein. The first obvious observation of the data is that (a) the peaks get smaller as the experiment proceeds and (b) the baseline is highly non-linear, despite the data itself being relatively noise free. In order to start analysing the data, we first need to deduce what is baseline and what is signal. Software provided with the instrument, as well as available as freeware written by experts, will correct for any baseline drift in the system. The data can also be offset by a finite amount due to the heat of dilution of the buffer components (a correction for this can also be applied). The amount of heat in each injection is represented by the area under the curve, and this is then presented versus the molar ratio of ligand/protein as the more familiar ITC thermogram data in Figure 2B: it is this data set that is then analysed. The first

quantity that can be extracted is the enthalpy of binding. This is the sum of the area under the curve of the original thermogram (or the sum of integrated heats from Figure 2B) divided by the molar concentration of binding sites. This gives a quantity in units of kJ/mol or kcal/mol. Errors in determination of the enthalpy of binding can arise from inaccurate representation of the baseline and as such, this baseline correction (and correction for any heat of dilution of the components) is of paramount importance if accurate thermodynamic quantities are to be derived.

The decrease in peak area as the titration proceeds is due to the decrease in available binding sites for the ligand as binding occurs. The degree to which this decline in peak area occurs will be due to the affinity of the ligand for the site. For very strong binding, there is effectively an 'all or nothing' effect whereby as each volume of ligand is titrated in it is bound by the protein. This situation is seen in Figure 3A where a set of large peaks is followed by a sudden transition to virtually no signal. For more modest affinities, some ligand is bound upon titration, while some remains free. This gives rise to the situation seen in Figure 3B, where there is a definite decrease in the peaks, and a pronounced curvature of the thermogram. At very low affinities, upon titration of the ligand, only a small portion of it binds to the protein. This decreases the size of the peaks and even at high ligand concentrations the titration doesn't appear to reach an 'end point'. Given what has been previously noted about the importance of a good definition of the baseline in order to accurately distinguish the heat of binding from the heat of dilution, it is clear that weaker interactions will be inherently more error prone in their affinity determinations. The curvature of the plot is therefore related to the strength of binding and as such at high affinities where there is little or no curvature, then it will be difficult to determine an affinity as there is little data in this region to fit. Thus, there

Figure 3. Titration of a protein with three different ligands. Each ligand binds to the protein with either **A**) a high affinity, **B**) a medium affinity or **C**) a low affinity. **Top panels:** experimental trace. The baseline is indicated in red. **Bottom panels:** the integrated heats upon injection (black squares) and the data fit (red line) after subtraction of the control data.



appears to be an 'affinity window' between interactions that are too tight and those that are too weak. Empirically, this has been found in ITC to be related to a value, C , defined by $C = K_a * M$, where K_a is the association constant and M is the protein concentration. This value should lie between 1 and 1,000, although values of 20–200 have been thought to be optimal. Methods for determining very tight and very weak interactions using ITC, which require multiple titrations, have now been developed (see 'Further reading' for more details).

To determine affinity, the data is fitted to the Wiseman isotherm, which derives the stoichiometry (n) and the affinity K_a . Once K_a is derived, using the relationship $\Delta G = -RT \ln K_a$ free energy can be derived, and then, using the 2nd Law of Thermodynamics ($\Delta G = \Delta H - T\Delta S$), entropy can also be determined.

For most standard 1:1 interactions of well-behaved proteins, where both ligand concentration and protein concentration are known accurately, this is where most analysis can stop. The thermodynamic parameters can then be interpreted in terms of known molecular interactions, and in many cases, related to molecular dynamics (MD) and Monte Carlo methods (MC) studies. Circumstances in which this is not the case will be considered in the next section.

Things that can go wrong or make your ITC thermogram look strange

The first thing that usually goes wrong in an ITC experiment is the value of n is wrong, or at least not 1, despite the fits to the data looking good. The origin of this is often an error made in the determination of concentration of either the protein or the ligand. This former reason will have several consequences. Firstly, as quantities such as enthalpy depend on the determination of the amount of

heat divided by the protein concentration, then inaccurate measures of protein concentration will lead to an inaccurate determination of enthalpy and consequently errors in the value of the other thermodynamic quantities. Secondly, n is part of the fitting routine for determining the affinity, thus if there is error in n , then there is a knock-on error in K_a , and consequently error in the free energy determination. Any error in n therefore leads to greater error in the derived thermodynamic parameters. Errors in ligand concentration will also have effects on K_a and are an additional source of error over and above that of the protein.

A second reason for strange-looking thermograms is that either the protein, the ligand or the protein ligand complex is prone to aggregation. Any aggregation will have its own heat signal which will be convoluted with the binding process. If this aggregation is to a definite state, i.e., ligand-induced dimerization, then there are fitting routines that can account for this. Where there is non-specific aggregation, then it is very difficult to assess the individual contribution from each of the aggregation processes from the binding process.

Ensuring that the ligand and protein are dissolved in the same buffer may seem trivial and is often overlooked by novice users of ITC. However, extensive experience has shown that this oversight can have a profound effect. Any difference in the two buffer solutions will give rise to what can be quite a substantial heat of dilution. Researchers should always dialyse their samples, pass them down a gel filtration column to achieve thorough buffer exchange or take them up in buffer from the solid state (making sure there are no dried contaminants). Many researchers try and convince themselves that such a situation won't occur for them, but 20 years of working with ITC has shown it to be a very unforgiving technique when it comes to buffer requirements. Dodging this step will almost always lead to uninterpretable results.

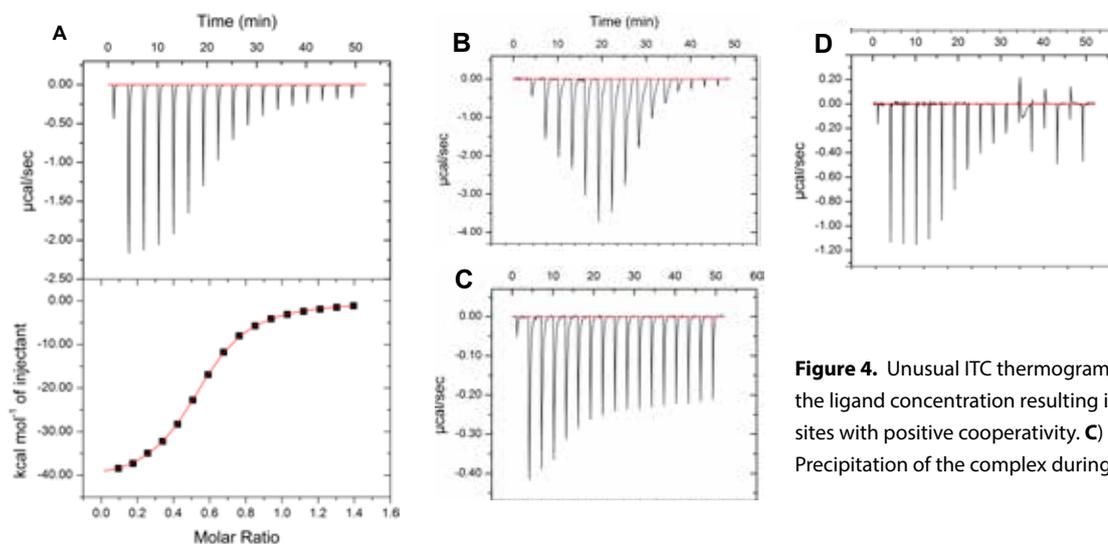


Figure 4. Unusual ITC thermograms. **A)** Incorrect quantification of the ligand concentration resulting in $n = 0.56$. **B)** Multiple binding sites with positive cooperativity. **C)** Poor buffer matching. **D)** Precipitation of the complex during the titration.

A third cause of a non-standard thermogram is the existence of multiple binding sites within the protein, with the potential for cooperativity between these sites in the form of allosteric interactions. Multiple sites may exhibit either positive or negative cooperativity, meaning that the binding at the first site strengthens or weakens the affinity for the ligand at the second site. Such thermograms tend to be biphasic in nature (e.g., Figure 4B) and it is possible to fit this data to models that generate values for affinity of the individual sites and for the thermodynamics of interaction between the sites. These can then be readily related to models of domain motion, where these exist, to provide a mechanistic interpretation of the thermodynamic signatures.

Finally, it is always wise to check whether your ligand or protein self-associate, as this will both generate an apparently anomalous heat signature and give erroneous values if fitting to simplified model of binding.

Conclusions

ITC has proven itself to be a powerful method for deriving thermodynamic data regarding binding and association processes. It has reintroduced basic thermodynamics into many biochemistry laboratories and enabled structural biologists to make informed interpretations of their data based on rigorously derived parameters. The majority of studies still feature a 1:1 binding event, and in this article, we have concentrated on how this is derived and what can go wrong experimentally. There are also situations involving multiple binding sites where ITC can yield important information about the interaction between binding sites, and the thermodynamic forces controlling these. New instrumentation and better analysis methods are now available putting the technique in the hands of more researchers. ■

Further reading

Basic ITC

- A good introductory protocol can be found online at: <https://www.nature.com/protocolexchange/protocols/2546>
- Claveria-Gimeno, R., Vega, S., Abian, O. and Velazquez-Campoy, A. (2017) A look at ligand binding thermodynamics in drug discovery. *Expert Opin. Drug Discov.* **12**, 363–377
- Baranauskienė, L., Kuo, T.C., Chen, W.Y. and Matulis, D. (2018) Isothermal titration calorimetry for characterization of recombinant proteins. *Curr. Opin. Biotechnol.* **55**, 9–15

More complex systems

- Vega, S., Abian, O. and Velazquez-Campoy, A. (2015) A unified framework based on the binding polynomial for characterizing biological systems by isothermal titration calorimetry. *Methods* **76**, 99–115
- Dutta, A.K., Rösgen, J. and Rajarathnam, K. (2015) Using isothermal titration calorimetry to determine thermodynamic parameters of protein–glycosaminoglycan interactions. *Methods Mol. Biol.* **1229**, 315–24
- Feig, A.L. (2007) Applications of isothermal titration calorimetry in RNA biochemistry and biophysics. *Biopolymers* **87**, 293–301



Dr Gemma Harris manages the Biophysical Instruments Facility at the Research Complex at Harwell, providing training and support for a variety of biophysical instruments. She completed her PhD, and subsequent postdoctoral position, at the University of York researching host–pathogen protein interactions before moving to her current position. Email: gemma.harris@rc-harwell.ac.uk



Dr David Scott is Reader in Physical Biochemistry at the University of Nottingham and Group Leader in Biophysical Methods at the Research Complex at Harwell. He did his degree in Molecular Biophysics at Leeds, followed by doctoral training. Subsequent postdocs were at York, Oxford and Bristol before appointment at Nottingham where he teaches first-year bioscience and MSc students basic chemistry and thermodynamic methods, respectively. Email: david.scott@nottingham.ac.uk