A non-technical summary of my thesis – Typical case behaviour of spin systems in random graph and composite ensembles

Jack Raymond
February 21, 2010

Introduction

The recent convergence of the sciences around common themes of complexity and emergence has given rise to new problems which cross the boundaries of traditional research areas. My thesis describes the behaviour of large collections of simple binary variables (e.g. bits of information) in three models representative of such multi-disciplinary research: a model of 'multi-user detection', an optimization problem called 'exact cover', and a 'composite model' that combining two model paradigms. Solutions to these problems impact on important open questions in mathematics, physics, computer science and other fields.

At a small scale the three problems tackled are distinct, as are the applications of their solutions. However, what links these problems is the common phenomena emerging when sufficiently many of the variables are studied collectively. It is possible to gain a better understanding of such problems by using randomness to describe disorderly features and by focusing on large scale properties. For this reason, the methods and techniques used in my research are drawn principally from the field of statistical physics, with input from information theory and computer science. An approach called the replica method is used extensively; this is a powerful tool that describes the model statistics in a simple framework. At the same time exact and approximate computational methods, called algorithms, are developed and analyzed. In this short essay I describe the main themes of my thesis, summarizing the key concepts and results at a non-technical level.

Multi-User Detection

Several sections of my thesis focus on a model of multi-user detection: wherein the signals, transmitted simultaneously to a base-station by many users, interfere with one another, and the information of each user must be extracted from the joint signal. In the case of mobile phones for example, user may each send a bit of information as a modulated signal on some bandwidth, the bandwidth being a set of frequency bands. Often the number of frequency bands is large, but in order to maximize the efficiency of information transmission many users will be packed onto the bandwidth, causing signal overlaps. When the joint signal is decoded at the base station an estimation of David's bit may not be possible without simultaneously making assumptions on the bits of Juan, Michael and Martin with whom he shared the bandwidth. It is said that the estimate for David's bit is constrained by other bit estimates.

In this example we are interested in a number of different properties. Some questions require an algorithmic approach, a computational method to go from a received signal and transmission model to an estimate of bits. For example, what is the best estimate of the bits given the signal, and how can this be constructed? At the same time some questions might be solved independently of an algorithm, for example: how much information can be reconstructed? At what level of interference does estimation cease to be effective? These questions are addressed principally with the replica method.

At the level of a few users we might answer all these questions in some exact and systematic manner, but if there are many users this may become impractical. For example if four users each transmit one bit of information, there are sixteen possible bit sequences to consider: we may evaluate a model probability for each one, and choose the most probable as an optimal estimate. However, throw in Ton, Laura, Jort and Gabrielle, plus a few hundred other users, and a solution by this method will require the age of the universe to solve on a fast computer, even for the simplest probabilistic model! This is because the number of possible bit sequences increases exponentially with the number of users. Unfortunately, in a 'worst case' of transmission, some particular set of bits and bandwidth access patterns, no method is guaranteed to give an optimal solution significantly quicker.

Despite the seemingly bleak outlook at the worst case level large problems are regularly encountered and solved efficiently by approximation, and in some cases methods have been shown to be (with high probability) optimal. In fact optimal outcomes are not uncommon, and the methods are often relatively simple. The reason is that the set of interfering user signals can be understood at a statistical level, with the worst case of interference a statistically insignificant outlier. For all practical purposes it is the properties of statistically plausible, so called 'typical cases', that are important. Excluding outliers, almost all cases tend to converge on a small, but descriptive set of common properties; in the same way that familiar concepts such as temperature and pressure describe accurately the collective phenomena of many molecules. In multi-user detection these properties can be exploited to understand the limitations of information transmission, and the development of fast and robust estimation methods.

The third chapter of my thesis deals with a Code Division Multiple Access (CDMA) model of multi-user detection. CDMA is a protocol used in wireless communication, and a special type of 'sparse' code was investigated where each user transmits only on a small number of the available frequency bands, in scenarios where the bandwidth is large. Transmission in this fashion has some advantages over the standard ‘dense’ implementation, where the full bandwidth is accessed by every user. It was found the performance indicators were comparable in most regards to the dense method, although some practical issues beyond the scope of the thesis would be central to any implementation. This work clarified the differences in large scale performance between the sparse and dense model forms.

Exact Cover

Another theme of the thesis revolves around an optimization problem called exact cover, which is a problem closely related to scheduling and a large family of challenging optimization problems. In exact cover a large set of tasks must be completed, and three candidates exist for each task. For example a list of tasks: A,B,C... ; Lenka may be capable of tasks 'A,B,D', Andrea may undertake 'B,F,G', Elitza 'B,E', Talya 'G,H' and so forth for many candidates and tasks. The solution to the following question is required: Is it possible to choose a subset of candidates to cover all tasks, without redundancy? An incomplete cover has too few candidates selected and some tasks go uncovered, whereas a redundant cover has some tasks covered by more than one selected candidate (e.g. if both Lenka and Andrea are selected the task B cover is redundant). An exact cover, if it exists, is a subset covering all tasks exactly once.

The theoretical issues involved in deciding if such a set exists, and providing a detailed solution, are closely related to multi-user detection, in fact one may identify special cases

where they are equivalent, making some results transferable.

The second chapter analyzes a "branch and bound" algorithm capable of discovering exact covers in large systems. What is found are two regimes, one where the algorithm always succeeds quickly, and one where it runs too slowly to be of practical use. The feasibility regime of this algorithm was compared to statistical properties of the model and solutions. This chapter contributes to a greater understanding of sufficient and necessary model features for computational efficiency in large scale problems.

Composite Model

The final problem is more abstract in nature, and is close to some standard models of statistical physics where energy and temperature are key concepts. Variables again take bit values, and are subject to pairwise constraints. The energy of a particular assignment is the sum of penalties incurred each time two coupled variables are of the same value (i.e. both 0, both 1). An understanding of the set of assignments at some energy level, i.e. at some temperature, is then sought. Low energy solutions are of most interest, but describing these is problematic due to energetic frustration - since in general not all pairs of variables can be made to simultaneously mis-align.

One common conceptualization of the model solutions involves 'energy landscapes'. The landscape elevation at a point (assignment) is determined by the energy of the assignment, related assignments (those with similar variable assignment) are positioned closely in the landscape, whereas strongly distinct assignments are distant in the landscape. Nearby points have comparable energies, so that valleys of low energy, and hills of high energy become apparent. In this context one can understand the set of assignments characteristic of some temperature as those touched by a ball bouncing through the landscape - as temperature is decreased the ball is less energetic and the configurations observed become restricted to valleys. These valleys are again described statistically, and it is the large-scale landscape structure that is of most interest.

With variations in the constraints, the landscape varies within three notable scenarios: one with no significant valleys, one with a single dominant valley, and one with many valleys. Two commonly studied model types are well known to show complicated but tractable behaviour. In the sparse model paradigm only a small number of partners constrain each variable; whereas in the dense paradigm every pair of variables is weakly mutually constraining. Considering scenarios in which both types of constraints are simultaneously present has until recently been a neglected topic. Although energy concepts are introduced in the context of this third problem, they are also central to my studies of multi-user detection and exact cover.

The fourth chapter introduces the model of mixed constraint types, the 'composite model'. Important landscape statistics were described with variation of constraint strengths and temperature, and a fast algorithm was developed to study specific examples. Composite models might be relevant in engineering, or as an approximation to models with different constraint types. The fifth chapter applies the new dense and sparse crossover to CDMA, each user transmitting with a fraction of power on a small subset of bands, with the remaining power spread thinly on the remaining bandwidth. Again, regimes of favorable performance was observed, emerging only as a large scale phenomena.

Conclusion

In conclusion, a number of problems were addressed using complementary techniques from statistical physics, information theory and computer science. This has shed light on the origins of large system complexity in the three problems addressed, and in related fields. New results for multi-user detection protocols were shown indicating feasible of new paradigms, a challenging optimization was demonstrated to be efficiently solvable when considering large instances, and the interplay between different constraint frameworks was explored.

My thesis work was undertaken between October 2005 and November 2008 with the Neural Computation Research Group at Aston University, under the supervision of Prof. David Saad. The work culminated in several research journal and conference publications. The thesis was submitted in November 2008 and successfully defended in January 2009, the examiners were Dr. Juan Neirotti of Aston University and Prof. Ton Coolen of Kings College London. I extend special thanks to my numerous colleagues who all provided more than a bit of inspiration.