Essay Review of *Quantum State Diffusion* by Ian Percival

Tim Spiller
Trusted E-Services Laboratory
HP Laboratories Bristol
HPL-2001-7
January 22nd, 2001*

E-mail: ts@hplb.hpl.hp.com

quantum state diffusion, measurement

Quantum state diffusion is a theory which has applications to both the foundations of quantum theory and to the behaviour of open quantum systems, quantum systems interacting with environments. This article is an extended essay review of *Quantum State Diffusion* — the first book [1] on the theory, written by one of its originators, Ian Percival.

* Internal Accession Date Only
Approved for External Publication

1 Introduction

The theories which come under the general heading of quantum state diffusion have impact at both ends of the spectrum of people interested in quantum physics. They provide tools for those interested in unravelling and modelling the evolution of open quantum systems—quantum systems coupled to environments—and understanding what goes on in quantum experiments. They also provide much food for thought for those interested in quantum foundations, particularly the measurement problem in conventional quantum mechanics and the quantum-classical boundary. This first book on quantum state diffusion [1], by Ian Percival (who has played a major role in the development of the whole field), has something in it for folk at both ends of the quantum spectrum, the pragmatic and the philosophical.

Part of the reason for this is that these folk are probably now rather closer together than they used to be, whether they like it or not! Over the last decade or so, state of the art experiments have been developed which probe individual quantum systems. Examples are atoms or ions (in beams or traps), photons (propagating or in cavities), electrons or Cooper pairs thereof (in tiny sub-micron condensed matter devices).[1] It is therefore no longer really possible for those who do or analyse such experiments to think just in terms of ensemble probabilities generated from some simple Schrödinger wavefunction—the decohering effects of other degrees of freedom and the results of measurements can be detected *event by event*. Quantum experimenters therefore (quite justifiably, in my opinion) like to ponder on the meaning and interpretation of such data. Similarly, it is therefore probably also unwise for those folk who like to spend their time trying to swat the measurement problem bug to ignore these experimental breakthroughs.

Indeed, it is input from both ends of the spectrum—stochastic modelling ideas for the description of individual quantum experiments and very careful thinking about just what is wrong with conventional quantum mechanics (at which John Bell was a master)—that has led to quantum state diffusion and related ideas. Percival stresses this. However, despite this, the book is constructed to permit omission of some later chapters, dependent upon your bent. Whilst you can certainly still get something valuable out of a partial read, I'd nevertheless encourage quantum experimenters to read about the foundations (even if you skip all the proofs) and philosophers to read about applications (even if you don't bother to dust off your PC and run a few simulations).

To a first approximation, the structure of this review follows that of the book.

2 Background

When I first learnt about quantum state diffusion, the general approach seemed entirely reasonable. In fact, more than that, it seemed to fill a gap. After all, in classical physics we

[1] Essentially, the full list is that of the current candidates for quantum computing hardware. Such hardware, if it is ever built on a large scale, will require the coherent control of a large number of such individual quantum systems.

have two rather different descriptions for systems interacting with environments. Suppose you chuck a load (an ensemble) of little particles into a big hot 'bath', where they get buffeted by thermal fluctuations. You can choose to follow one of the particles. It does a random walk—it jiggles around stochastically, exhibiting Brownian motion. Alternatively, you can choose to ask what happens to the whole ensemble. The average density diffuses smoothly. These two descriptions are rather different. Nevertheless, it's possible to average over the jiggling individuals and show that, indeed, the ensemble exhibits nice smooth diffusion. If you want, you can therefore get picky about the name, because in the quantum case a 'diffusing' quantum state jiggles around, describing the stochastic evolution of an individual member of the ensemble—the analogue of a single Brownian particle. The average over these fluctuating states gives the smooth ensemble behaviour. The latter is the familiar density operator evolution, the way we are all taught (sometimes with the implication that it is the only way) to handle quantum systems talking to environments.

To warm up, the book kicks off with the classical Brownian motion example. This is nice because it's easy, it reminds you that there is more than one way to treat an open system and it helps to introduce vital bits of probability and statistical luggage and the Ito calculus. "The 'what' calculus?", some of you may well ask. The Ito calculus is a nifty differential calculus for handling stochastic fluctuations. It looks a little strange if you haven't seen it before, because the mean of a noise differential squared is a time differential. However, this enables the handling of the mean displacement squared growing linearly with time in a random walk. It also makes it clear that for systems subject to both drift and such noise, the drift dominates for large times and the noise for small times (in units of an identifiable characteristic time). If you plan to work through the derivations later in the book, it's probably sensible to make sure that you are happy with all the calculations discussed in the context of Brownian motion. Just to re-emphasise, there is nothing mysterious and quantum about the Ito calculus, so it's worth making friends with it in a simple classical situation.

3 Open Quantum Systems

An open quantum system is one coupled to other stuff. This other stuff—the environment—could just be the other half of an EPR pair, or it could be a moderately complicated quantum system, or it could be a huge bath of quantum degrees of freedom at finite temperature. Even if there is good reason for assuming that the whole system—system of interest plus environment—is well isolated from any further stuff and so evolves according to the good old Schrödinger equation, neither the system nor the environment viewed alone will. In general, if they interact they become entangled. If you have no infomation about what the environment is up to, or if it so unwieldy that you can't really calculate what it is up to, then, even just as a practical tool, it is handy to have a description for the evolution of just the system of interest, the open quantum system. The familiar way to do this is to define a reduced density operator for the open system, 'reduced' because the environment degrees of freedom have been traced over, in effect averaging over the entanglement between system and environment. The resultant so-called master equation evolution is non-Schrödinger, although still linear. This approach is fine for calculating the average behaviour of an ensemble of systems (although it may not be quite so fine in practice if this requires solving a whacking great matrix evolution equation). However, it is certainly not so fine if you want some idea about what the individual members of the ensemble are up to.

Of course, some very eminent quantum folk (including Neils Bohr and Rudolph Peierls, I suppose) held the opinion that this question shouldn't really be asked in the first place. This view still has some support. Me? Well, I look at an experimental plot of the jumps between fluorescence and no fluorescence generated by a single ion in a trap (due to its shelving in a metastable state), and I think it's a question well worth asking. Especially when you chuck in the fact that it is possible to explain and picture classical diffusion of open systems at the individual level. Convinced yet?

The natural description of an individual open quantum system is by a quantum state. When there is no environmental coupling this clearly evolves *à la* Schrödinger. However, to allow for—to model—the effects of entanglement with an environment, the evolution must be non-linear. (It turns out that it also has to be stochastic; this is discussed in detail later on in the book.) A quantum measurement can clearly be regarded as a certain sort of interaction with an environment, so it isn't singled out as special in this approach.

All of this is discussed in the build-up to the introduction of the quantum state diffusion model (QSD). Another important point which is made concerns the ambiguity of the boundary between system and environment. This boundary is clearly ficticious. You could put it in one place, but your mate describing the same overall system could choose to lump another quantum degree of freedom (or two) into her/his definition of the system of interest, even if just for the hell of it. Both descriptions are permissible (and they had better agree when used to calculate the same physical quantities). It is therefore clear that using a jiggling quantum state (in QSD, or, for that matter, any other approach) to model the effect of entanglement with an actual environment cannot be fundamental. QSD doesn't solve the measurement problem, nor does it claim to. However, that doesn't mean that it isn't useful and doesn't provide some very valuable insight and input for future work. The pictures produced by QSD are very interesting, particularly those modelling measurement; however, they shouldn't seduce you too much.

4 Quantum State Diffusion (QSD)

Following all the above background, chapter 4 introduces the quantum state diffusion model (QSD). The starting point is the Markovian form of the master equation, due to Lindblad. In this, the Schrödinger evolution of the reduced density operator for the system of interest is modified by additional terms. These contain operators (so-called Lindblads) which allow for the effects of interaction with an environment, generating non-unitary evolution of the density operator. To produce QSD, an evolution equation is derived for a quantum state (representing an individual system) so that the ensemble average of this state evolution generates the master equation, analogous to the average over Brownian trajectories giving smooth ensemble diffusion in the classical case. QSD is both non-linear and stochastic, so the ensemble average means an average over the noise. This is where the Ito calculus comes in.

QSD can look a little scary at first, but examples always help. Lindblads can be Hermitian or not—in general Hermitian Lindblads model measurement interactions and non-Hermitian Lindblads model decay or dissipation. Measurement is illustrated with example Lindblads such as a spin component of a spin-1/2 system and the number operator (energy) of an oscillator mode. Decay is illustrated with the spin lowering operator and oscillator mode annihilation operator. It is clear from these simple examples that the individual trajectories exhibit stochastic jiggling, but over time they demonstrate localization, such as converging to an eigenstate of an operator being measured. This localization is very important and it is discussed in detail later on in the book.

The construction of a state evolution which on average reproduces a particular master equation is often referred to as an unravelling of that master equation. Although it is derived following certain assumptions, QSD is not the only possible unravelling of the Markovian master equation. This is pretty obvious really, since it is clear that in general a given density operator does not have a unique decomposition into an average over pure states. Many different ensembles can be responsible for the same density operator, so there are different possible unravellings. For example, it is possible instead to work with quantum jumps—finite discontinuous changes of state instead of the differential stochastic changes of QSD—and still unravel the Markovian master equation. For the case of Hermitian Lindblads, it is possible to work with real differential noise instead of the complex noise of QSD. All this is not a terminal problem for QSD, or any of the other approaches—indeed, it is inevitable given that there is a one-to-many relationship between density operators and ensembles. This is a further illustration that QSD and the other unravellings are not fundamental theories. They all model the apparent loss of a system's coherence (due to its entanglement with an environment) by actually destroying it. This begs a question as to the choice of unravelling. Although there isn't an answer, sometimes further information about the nature of the environment can contribute to the decision-making. For example, in the case of actual measurement (say of the z-component of a spin) it would be appropriate to use QSD, whereas if the actual environment is a noisy Hamiltonian term (magnetic noise in the z-direction) then it is more appropriate to use an unravelling with real fluctuations. These two situations have the same spin master equation and so the same evolving density operator, but this arises from rather different ensembles in the two cases. Without any additional information, the unravelling choice may simply come down to computational convenience and efficiency. There is some discussion of alternative unravellings in the book, with references of course. However, beginners might prefer a little more to hand,[2] say with some actual examples calculated with quantum jumps for comparison.

It's worth noting a couple of nice developments which have emerged since the book was produced. A QSD model has now been constructed to handle non-Markovian evolution [2, 3, 4, 5]. This work shows that the idea of unravelling is not simply restricted to the Markovian case and it extends the usefulness and applicability of QSD. From the simulation perspective, the non-Markov extension of QSD is particularly useful in a truncated perturbative form [4]. Detailed discussion, with examples, has also been given on the position of the system–

[2]Although, of course, this can be tracked down through the references...

environment boundary, showing how the approach deals with (and is consistent with) different choices of the boundary position [2, 3].

5 Localization and the Classical Limit

A nice feature of QSD is that in general it produces localization of a quantum state as it evolves in time. This is in keeping with what we say happens during quantum measurements—the system is projected to an eigenstate—and is also consistent with our everyday perception of classical systems following well-defined trajectories. Various aspects of localization and classical behaviour are discussed in three chapters of the book, 5, 9 and 10. These are rather more mathematical than the rest, but it's worth taking on board their content even if you skip the proofs.

Clearly if QSD is to simulate quantum measurements, it has to reproduce the Born rule. That is, if a pure state which is a superposition of eigenstates of the quantity to be measured is subject to an appropriate measurement interaction, the outcomes should converge to the eigenstates with a probability distribution given by the initial amplitudes (squared). Of course, this works. Intuition on the convergence to an eigenstate can be obtained by direct inspection of the QSD equation—it is clear that the fluctuation term vanishes when the system is in an eigenstate of the acting Lindblad. Detailed discussion of such localization is given in chapter 5, including analysis of the rate.

It's worth noting that the recent work on the system–environment boundary position [2, 3] has particular relevance to the measurement case. Clearly a better model of a measurement is really the convergence of the system state to an eigenstate along with the sympathetic convergence of some other system (the 'pointer' of the apparatus) to a state which enables identification of the measurement result (the eigenvalue), all this following the Born rule. Any description of a measurement interaction, such as QSD, clearly has to facilitate movement of the boundary, so that if part of the original system environment is defined to be a new (pointer) system, everything still works. Obviously this doesn't remove the ambiguity of the boundary position, but such consistency with different choices of the position is clearly necessary and it is about as good as you can hope get with a non-fundamental description of measurement.

Localization also happens for non-Hermitian Lindblads representing, for example, a dissipative environment. With continuous variables, such localization happens in phase space. It is also not complete—there is no violation of the uncertainty principle and so even localized states demonstrate the uncertainty relationship between conjugate variables. An example is the damped driven harmonic oscillator. Essentially you can start it wherever you like and after a while the evolution will converge to a localized state following the expected classical trajectory for such an oscillator. At zero temperature the width of the state is set by quantum fluctuations (thus maintaining quantum uncertainty), but for a finite temperature environment the width of the state is set by the temperature. All this is very appealing, as it shows how individual quantum motion can go over to exhibiting classical behaviour—a localized lump of state following a classical path—when the conditions are right. Clearly, as the motion in phase space becomes very large compared to $\hbar$, more and more quantum states become involved. (Quantum calculations become harder, although there are some nice tricks to help such as using a moving basis, discussed under 'Applications'.) Nevertheless, it is possible to get a feel for the emergence of classical behaviour from such considerations.

More detailed and formal discussion of the classical limit is given in chapters 9 and 10. In a sense, this comes at it from the other direction, considering quantum corrections to classical motion. The analysis is done for continuous variables by linking a classical density of points in phase space to a quantum density (which doesn't allow for quantum wave phenomena but can handle localization) and through to the Wigner function (which does include quantum wave effects). The general semi-classical limit of QSD is not simple to analyse. Handwavingly and ignoring finite temperature, when a wave packet is well spread in phase space compared to $\hbar$, localization bites until the packet shrinks and follows an approximately classical path—as seen in numerical examples. As always, it is easier to analyse the example of linear systems and Gaussian wave packets. The evolution of an individual system is described by a Gaussian, whose path is classical but whose fluctuating shape is quantum. If you don't look too closely, this is a classical particle! Most of the formal detailed discussion of localization is probably for aficionados; however, it's worth appreciating that such a classical limit can be taken.

It is clear that localization and the semi-classical motion know all about a frame. For example, with dissipation a packet can essentially come to rest. However, this is no more strange and a headache for relativity than friction is in classical mechanics, and the answer is just the same. The frame is singled out by the environment. Just another reminder that QSD isn't fundamental.

6 Applications

Chapter 6 is devoted to some example applications and a lesson on how to write a QSD program. The latter is possible even if you don't want to get your hands (very) dirty, because Rüdiger Schack and Todd Brun have done all the hard work for you [6]. You can even get the code for free over the internet from: http://www.ma.rhbnc.ac.uk/applied/QSD.html . From the computational perspective, it is rather easier to solve for the evolution of a quantum state in QSD compared to the whole density operator—a vector takes less RAM than a matrix. The trade-off, of course, is that many QSD runs are needed for decent statistics. However, sometimes it's possible to get half-decent statistics from a relatively small number of runs, and quite often it's possible to learn a good deal about the system from just a few runs. Another handy trick (which rarely helps in a numerical attack on the density operator) is to utilize the localization inherent in QSD to significantly chop down the basis of states actually needed to represent an evolving state.

One of the nicest applications of QSD to date is the modelling of dissipative quantum chaos.[3] These studies have shown how individual QSD wave packets localize and follow a classical path, even when that classical path is chaotic. Cute pictures! Other examples discussed in chapter 6 are second harmonic generation, Stern-Gerlach measurements and

[3]Okay, okay, I'm biased as Jason Ralph and I were the first to consider this...!

noise in quantum registers,[4] showing the wide range of applicability. Of course, this range has further expanded with the non-Markovian dvelopments [4].

7 Foundations of Quantum Theory

I've stressed more than once that QSD isn't fundamental. Nevertheless, and in addition to its applications, QSD does exhibit features worth retaining in a new fundamental theory. This does require you to buy one thing: the reality of matter waves. Ian Percival is obviously convinced about this and I'm pretty sure that I am too. If you're not, then you probably won't go along with the subsequent reasoning. However, in this case it's probably worth just stopping to think how instead you expect the problems of quantum mechanics to get resolved. Maybe you don't care? After all, as John Bell reminded us, quantum mechanics works FAPP (for all practical purposes). This isn't good enough for Percival, or me.

If matter waves are real in the same way that familiar classical things—made of matter—are real, then it's clearly sensible to try and cover the lot in a single unifying theory. Matter should be wave-like when it's quantum and lump itself into things looking like particles when it's classical, and there shouldn't be any need for observers and all that stuff. QSD incorporates both quantum wave behaviour and localization, but as it only models the effect of an environment and it blows a raspberry at relativity, it clearly won't do.

All this is discussed in chapter 7, with brief mention of existing alternative theories to quantum mechanics, in the build-up to...

8 Primary State Diffusion

The idea is very simple. If the diffusion of quantum states isn't just a model for the effect on them by other quantum systems (so that, in principle, the whole lot—system plus environment—could be treated quantum mechanically), then instead the diffusion must be *fundamental*. Everything else should follow, including, where appropriate, accurate Schrödinger evolution for isolated systems. In chapter 9 Ian Percival has a couple of cracks at constructing such a Primary State Diffusion (PSD) theory.

The first clearly still blows a raspberry at relativity, and so it isn't presented as the answer. Nevertheless, it has some interesting features. Starting with just diffusion (that's right, no Schrödinger evolution), which will dominate at short times, a longer term PSD evolution can be found which looks like Schrödinger wave behaviour, but with additional terms (just like in QSD) which cause localization in energy. Clearly not the ultimate theory,

[4] As something of an aside, it's worth noting that quantum state diffusion has potential relevance for quantum computing at both the pragmatic and the foundational ends of the spectrum. On the one hand, it is a useful tool for modelling and understanding decoherence in quantum gates and registers. On the other, if QSD points the way to an ultimate theory which is *not* quantum mechanics, then future attempts at quantum computing with many thousands or more of quantum bits could fail not because of environmental decoherence, but because of fundamental deviations from Schrödinger evolution for such systems.

because of relativity and because it doesn't fundamentally localize spatially extended systems (unlike, for example, the theory of GianCarlo Ghirardi, Alberto Rimini and Tullio Weber [7]). However, it's interesting that it can be done.

The second attempt takes a different approach, borrowing from the thoughts of Richard Feynman, Roger Penrose and others that gravity is really to blame (or thank?!) for collapsing wavefunctions. The starting point is now a Schrödinger-like evolution, but with spacetime fluctuations. Again, a PSD equation results. The problem is that there is no unique choice for the geometry of the spacetime fluctuations (to be consistent with relativity). You may think all this is a million miles from any experiments but, surprisingly, this is not so. Even with the Planck time (a miserly 5×10^{-44} s) setting the scale of time fluctuations for non-relativistic systems, existing matter interferometry results (i.e. the fact that interference is seen) actually rule out some possible forms of spacetime fluctuations. Of course, it will be extremely hard to actually look for evidence of PSD (or alternatives) in future quantum experiments. Any *fundamental* deviation from Schrödinger behaviour will have to be clearly distinguished from more mundane and almost certainly dominant environment interactions.

9 Comments

I had a considerable advantage reviewing this book. I read many of the papers on quantum state diffusion before I read the book.[5] Consequently, the book pulled together very nicely all these separate papers into a logical progression.

I also had a considerable disadvantage reviewing this book. I read many of the papers on quantum state diffusion before I read the book! This means I have difficulty in judging just how it would have read had I not done this.

I do very much like the book and in my best estimation it can be read 'cold' (i.e. without having read founding papers first). It is relatively short. This is a real plus because you don't have to wade through loads of excess baggage to find out about quantum state diffusion (even to get as far as being able to use it). It is also a bit of a minus because I do think a little more material in the odd place would further enhance the book. For example, I think more discussion on the different master equation unravellings (of which QSD is one)—maybe with some pictures comparing QSD with quantum jumps—would help folk new to the subject. It might also have been nice to include some real experimental pictures (such as a single ion in a trap jumping between states), to motivate and strengthen the whole case for a treatment of *individual* quantum systems. I was already a convert on this before reading, but others may need more persuasion.

Reading the book (or at least most of it) is very akin to listening to an Ian Percival lecture. Very thoughtful, very structured and fairly economical, but with repetition of key points just to make sure you've got the message. (The last couple of chapters contain more technical detail and are probably closer to the relevant original papers than the other chapters.) I think it is nice to finish off reading the book with chapters 7 and 8, as these are the

[5] I even wrote some of them on various applications...

pointers to future fundamental work. I therefore suggest that the final technical chapters, 9 and 10, are read (probably only by aficionados) after 4 and 5, or are regarded as appendices.

All in all, I think this is a very good book and I strongly recommend it. There is something in it for those who model and simulate open quantum systems through to those who continue to ponder about the foundations of quantum mechanics, and so it is particularly appealing to those of us with interests at both ends of this spectrum.

References

[1] I. C. Percival, *Quantum State Diffusion*, (Cambridge University Press, 1998, ISBN 0 521 62007 4).

[2] L. Diósi, N. Gisin and W. T. Strunz, “Non-Markovian Quantum State Diffusion”, quant-ph/9803062.

[3] W. T. Strunz, L. Diósi and N. Gisin, “Open system dynamics with non-Markovian quantum trajectories”, *Phy. Rev. Lett.* **82**, 1801 (1999), quant-ph/9803079.

[4] T. Yu, L. Diósi, N. Gisin and W. T. Strunz, “Non-Markovian quantum state diffusion: Perturbation approach”, *Phys. Rev.* **A60**, 91 (1999), quant-ph/9902043.

[5] W. T. Strunz, L. Diósi, N. Gisin and T. Yu, “Quantum trajectories for Brownian motion”, quant-ph/9907100.

[6] R. Schack and T. A. Brun, “A C++ library using quantum trajectories to solve quantum master equations”, *Comp. Phys. Comm.* **102**, 210 (1997), quant-ph/9608004.

[7] G.-C. Ghirardi, A. Rimini and T. Weber, “Unified dynamics for microscopic and macroscopic systems”, *Phys. Rev.* **D34**, 470 (1986).