Complexity and Analogy in Science
Pontifical Academy of Sciences, *Acta* 22, Vatican City 2014
www.pas.va/content/dam/accademia/pdf/acta22/acta22-phillips.pdf

Quantum Simulation of Calculationally Intractable Problems and Questions About Emergence[1]

William D. Phillips

I am from a place called the Joint Quantum Institute. This is a new institute that was set up just a few years ago to study quantum coherence.

A number of people at this meeting have quoted from the prologue to the programme of this meeting, which I think is an indication of how good that prologue was in guiding our thinking. The quotation that I wish to highlight is really a question:

> "Is there a difference between complex and complicated, such that some complex systems are not actually complicated even though all complicated systems are indeed complex?"
> (Werner Arber and Jürgen Mittelstraß, Prologue to the Program of the Plenary Session)

I intend to answer the first part of this question in the affirmative, from the perspective of a quantum physicist. I also want to say that I am not at all sure about the second part of this question. That is, there is an assumption made that all complicated systems are indeed complex. I'm not sure that that's true – but I'm not really prepared to discuss that yet. I think it might be in an interesting thing to discuss later.

Part of the problem is that we've been struggling with the definitions of what we mean by complex and complicated, and even simple. So I'm going to give you my definitions and they're the ones that I'm using for this talk. I'm saying that a problem is complicated if it's hard to state what the problem is, if it takes a lot of description to tell you what the problem is. My definition of complexity is that a problem is *complex* if it is difficult to solve. I think this is very much in line with the spirit of Kolmogorov Complexity.

[1] This text is a verbatim transcript of the talk given by William D. Phillips at the November 2012 Plenary Session of the Academy. The transcript has been slightly edited for clarity, and figures of slides to which the transcript refers have been inserted.

So *complicated* means that it is hard to state and *complex* means that it is hard to solve. This may not be the definition that everyone is using, but it is the one I'm going to use for this talk. An example of a problem that is both complicated and complex is the human brain. It would be very difficult to describe the problem in the first place, to describe all the neurons, and even more difficult to describe the initial conditions, and then it would be essentially impossible, even given all that description, to predict what happens next.

On the other hand, a pendulum is a very simple thing. You typically only need one parameter to tell you about the pendulum: the length of its suspension. Then, if you know the initial position and velocity, you know how the pendulum behaves after that. There is no chaos, everything behaves in a very nice way. So these are two extremes, a problem that is both complicated and complex, and a problem that is simple and *not* complex.

An example of what I consider to be an uncomplicated problem that is nevertheless complex is the Hubbard Model of Condensed Matter Physics. The idea is that you have a certain number of lattice sites. These might be the lattice sites on a crystal. And you have a certain number of particles. These might be electrons that can move through the crystal. In this model you only think about two physical parameters: one is the rate at which a particle can tunnel, that is, hop from one site to the next, and the other one is, what is the energy cost for putting two particles on the same site? That is, if I have two electrons, and they are on the same site, there is a Coulomb repulsion between them, and that costs a certain amount of energy. So the hopping makes things want to be everywhere throughout the crystal and the on-site energy cost wants to spread things out so that they stay in their own sites. Now this is a very simple problem to state, but if you wanted to calculate what happened, even at zero temperature, and you had more than a few tens of lattice sites and particles, no computer in existence can calculate this problem directly. And if the problem were even slightly larger than a few tens, then no computer that you could even imagine would be capable of calculating this problem. Why is that the case?

Consider a somewhat simpler situation. Imagine that I have N sites and on each one of the N sites is an atom, and that atom can be in one of two different states, and I'm going to call those two states '0' and '1'. If I've only got one site, then there are just two possibilities: the site is either 0 or 1. If I have two sites then there are four possibilities: I can have 00, 01, 10 or 11. If there are three sites then there are eight possibilities. So you see the number of possibilities grows exponentially with the number of sites. If I have just 300 sites, which is a very modest number, then that means I have 2^{300} different states. And that number is larger than the number of particles in the universe.

This means that a direct calculation for a quantum problem like this, would be really hard. Quantum mechanically the system can be in ALL of those states simultaneously. So that means there could be 2^{300} states simultaneously describing this system, and obviously no computer could handle that calculation, because the computer would have to be larger than the visible universe. So, it is fundamentally impossible to calculate a problem like this.

Why do we care? There is an important problem that is like the first one that I told you: the problem of electrons on lattice sites. One of the reasons why this is an important problem is that some people believe that this problem can describe high temperature superconductivity. Now this is not the ordinary kind of conductivity that, for example, Professor Zichichi mentioned in his talk. Superconductivity means the complete absence of electrical resistivity. Certain kinds of materials become superconducting at quite high temperatures, that is, temperatures on the order of a 100 degrees above absolute zero, whereas most ordinary superconductors are superconducting at a few degrees above absolute zero.

Such a "high temperature" is still very cold by the standards of daily human life, but very high compared to the temperatures of ordinary superconductors. And no one understands what causes this high temperature superconductivity. But some people guess that this very simple model that I described to you, will explain it. The problem is that even though the model is very simple, it is so complex that no one can calculate whether this model leads to superconductivity. So here is a model that is very easy to state, but the computational complexity of solving this model is so great that we do not know whether this answers the question of why we have high temperature superconductivity. And it's a very important question because being able to answer that question could help us to develop these materials for practical uses, which could be very important.

So the question is: is there any way out of this problem? And the answer is: yes, maybe (see Fig. 1). There are some approximation methods, but none of them work well for this problem. So that's not the way this is going to work, at least not for the moment. That is, calculational approximations so far have not been able to give us the answer to this problem. But there *is* a possible answer, and that is: doing experiments. It is something that I am particularly interested in because in my own laboratory, we can do experiments to, in fact, solve the problem that you see in Figure 1. This is one where the particles are not electrons, which are more complicated because they are fermions (they obey Pauli's Exclusion Principle, and I can't have two in the same state on the same site, so that makes things more complex as it turns out).

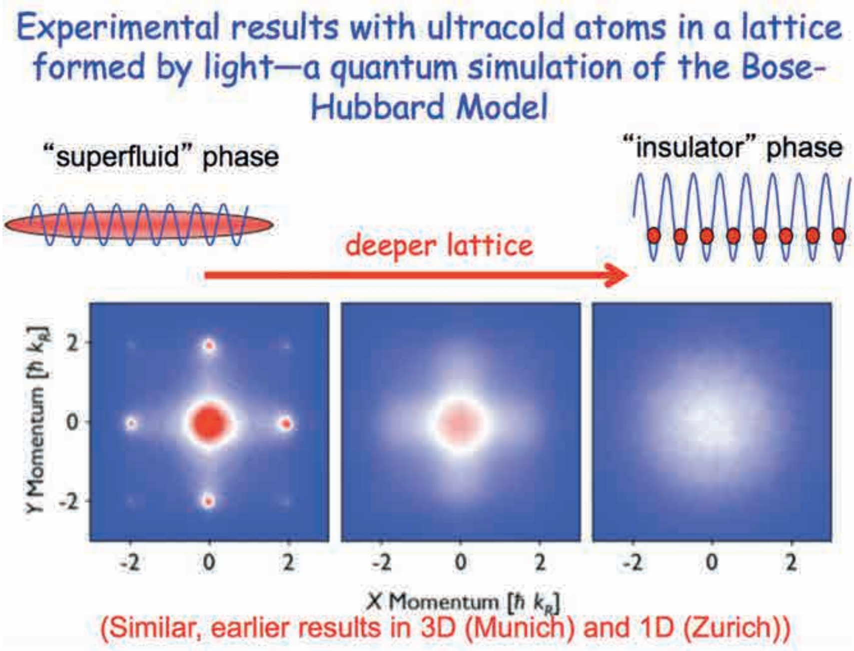

Figure 1. The Bose-Hubbard model.

This is the problem of what are called 'bosons', of which there are a number of examples (and the Higgs boson is perhaps the most famous one today) but the ones we use in the laboratory are things like rubidium atoms, which are bosons. It's an easier problem. We can solve the problem by approximation methods and we can solve the problem in the laboratory by making, in the laboratory, a physical realisation of the model. So it's not that we make something that is the same as some solid-state system that we think might be described by this model. We make the actual model in the laboratory. And what we find is that this laboratory model gives the same result as these approximation methods. So this gives us confidence that these approximation methods, in fact, work.

Figure 2 shows some of the results from our laboratory. I won't explain what these are because we don't have enough time. This is just to show that we *can* make measurements in the laboratory. And, by the way, I should say that the measurements that we've made in the laboratory have come after some pioneering measurements made in Theodor Hänsch's lab in

Munich, which is noted at the bottom of the slide. But the point is that in the laboratory we can make a physical realisation of this simple model, and we can show that it agrees with some of the most advanced approximation methods for the problem that we cannot solve directly, by brute force. Now the important thing is that we may be able to do the slightly more complex problem, the one with fermions instead of bosons, the one for which no one has solutions – we can also do that in the laboratory, or at least we hope that in the coming few years we will be able to make this model in the laboratory.

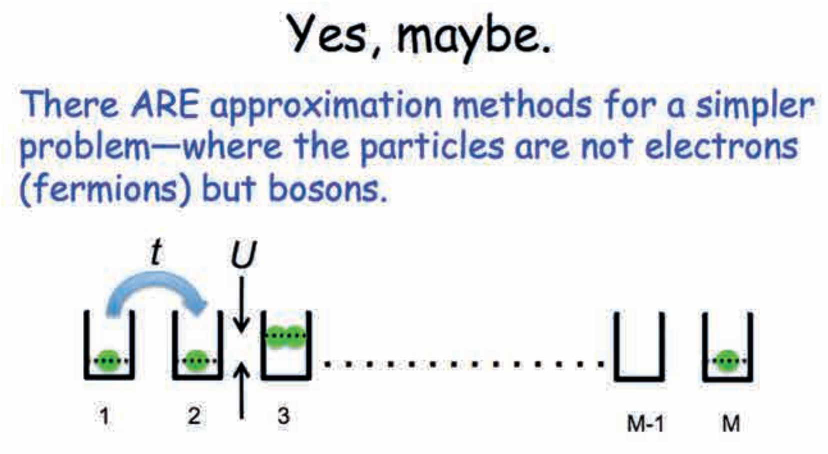

Figure 2. Experimental results for a laboratory realization of the Bose-Hubbard model.

What this means is that we have a problem that is so complex that, computationally, no one can solve the problem, no one expects to ever be able to solve the problem directly, and no one has yet been able to come up with an approximation method to solve the problem mathematically. And we believe that we can solve the problem in the laboratory by building a laboratory model of the mathematical model, and find out what the solution is we don't know.

What if it turns out that this model *does* show that there is superconductivity? Well then many people would say that what we have shown is that superconductivity, high temperature superconductivity, is an emergent phenomena, and I think this would fall under the category of what Jürgen

Mittelstraß called 'strong emergence', where if there is no possibility of calculating what the result is, but we know what the result is, in this case because we did an experiment, but in a certain sense that experiment is a kind of calculation.

So we might ask the question: what is emergence? And the answer is: I don't know! But I've been reading lots of definitions and some of them seem to be more philosophical than scientific. And, in fact, I was asked to give a talk at a meeting, in which the topic was 'Emergence at the mesoscale', and I complained to the organisers that I didn't know anything about emergence and I didn't know anything about the mesoscale. And they said, "Oh it doesn't matter, come anyway and just talk about whatever you know". So I did and I began my talk by telling people this: atomic physicists are confirmed reductionists. We are wedded to the belief that we can understand the beauty of nature on the basis of microscopic properties and principles. To us, emergence is just another name for *ignorance*. And I added a little winking emoticon to try to soften the harshness of this statement, but still I believe, as an atomic physicist, that emergence is just another name for ignorance.

So what is an emergent property? Here is my naïve attempt at a few definitions: it's a property of an ensemble, a macroscopic property that is not evident from the behaviour of the individual parts. Now, there are two possibilities, and I think that these two possibilities correspond to what Jürgen has called weak and strong emergence. One is that it's a macroscopic property that is not evident, and that no one would have been able to calculate beforehand from the microscopic behaviour, but that once that macroscopic behaviour becomes evident then you can figure out how to calculate it from the microscopic properties. That's the first possibility.

The second possibility is an ensemble property that *cannot* be calculated from the microscopic properties, even once the ensemble property, the macroscopic property, has been discovered. And that would apparently be the situation with the Fermi-Hubbard model and high temperature superconductivity.

So now I pose the following question: what happens if an incalculable problem becomes calculable because I change the hardware? In other words, I have now developed a new kind of computer. In my laboratory the computer is a bunch of atoms in a vacuum, and these atoms do the computation that no ordinary computer can do. So if a complex problem, one that was incalculable, all of a sudden were to become simple, or at least tractable, because of this new kind of computer, then how would this change our notions of emergence?

Consider another candidate for emergence, one that I have brought up before: human consciousness, that is, the awareness of self and all that that implies. How does that arise from chemistry and physics? I don't think anyone knows, and I suspect that no one will know anytime soon. Some people say that consciousness is an emergent phenomenon. Surely, if it is, it falls into the second category of being strongly emergent. So I offer that as a candidate for strong emergence.

Now in reading about emergence I found this quotation from Mark Bedau, who is a philosopher. He said that, "Although strong emergence is logically possible, it is uncomfortably like magic".

In the case of human consciousness, let's substitute the word "transcendence" for "magic". I am a person of faith and I have often been attracted to the notion that consciousness – whatever that really means, and I'm not sure I know – as well as free will – whatever that really means, and I don't think I know either – are transcendent phenomena. That is, they are gifts from God that are given to a sufficiently complex structure, maybe not just to humans. Is this notion any different from the idea of emergence? Is the idea that consciousness is a transcendent phenomenon, that is, something that has to do with the divine, any different from the idea of emergence, and is it any different from the idea of magic, as was indicated by Mark Bedau?

And then finally, is it any different from what I would call the old fallacy of ascribing to God everything that you cannot explain? This is sometimes called the 'God of the gaps'. Would the notion of transcendence or emergence change if we had a sufficiently powerful computational method like quantum simulation or quantum computation?

So why am I posing all these silly questions? It's simple. It's because I want to know the answers. I don't know the answers to any of these questions, and I'm not even sure that I even understand what the questions mean. And I excuse myself by saying: well, I think I'm a reasonably competent physicist; I'm certainly a bad philosopher, and an incompetent theologian. So I hope that a discussion of these questions will occur here in the Academy, and I welcome that discussion, because I hope that it will help me to become better educated.

Finally, I want to acknowledge the people that I work with, especially Gretchen Campbell, Ian Spielman, Paul Lett and Trey Porto, who are the permanent members of my group. They are all working together to try to provide the physical means of calculating problems that are impossible to calculate on computers.

I now want to invite your questions. The fact is that this talk is designed not so much to invite questions from you as to invite answers from you.

Let me summarize the questions to which I want answers, starting with a question that I didn't address: can a complicated system be other than complex? What is emergence – and this I think is strong emergence – and is it any different from magic? Does the idea of emergence change if you have new computational tools that change the tractability of a problem? Is there any validity in connecting emergence to transcendence? And does making such a connection bring one back to the old trap of ascribing to God the things that one doesn't understand. Normally I wouldn't be raising these last two questions. But in an Academy where some of the members are philosophers and theologians, I feel justified in raising such questions.

Thank you very much.