Photons

Newton thought that light consists of point particles, called corpuscles. Huygens, a contemporary of Newton thought that light is a wave. Not before Young did his diffraction experiments with single and double slits was this dispute temporarily settled in favor of the wave theory of light. In fact it took more than a century to realize that light is not a continuous wave, but in fact sometimes behaves in ways better described with the help of point particles. However, light particles, or photons, are unlike anything Newton ever dreamt of: they are massless, have spin, and, in a vacuum, move at the highest possible speed allowed by physics. In addition, photons may be polarized, a property that cannot be explained on the basis of classical point particles. Now that we are back to thinking of light in terms of particles, we again encounter the difficulties faced by the old corpuscular theories of light, e.g., the difficulty of explaining diffraction and interference. Quantum mechanics resolves this by assigning both wave-like and particle-like aspects to photons (i.e., the wave-particle duality).

Wave-Particle Duality

Because of their mutually exclusive properties all phenomena we encounter in our macroscopic, classical world can be grouped naturally into two categories: particles and waves. This natural division of the classical world does not hold in the microscopic quantum world: The objects of the micro-world exhibit both wave and particle aspects. Even the simplest quantum experiments, such as the single-particle double-slit experiment, cannot be understood without invoking this duality. What is this duality? Does it mean that photons, electrons, and other fundamental objects of our world are "really" particles if we look closely enough, and that the wave description is only a convenient approximation whenever many particles are involved simultaneously? In the case of light this idea seems natural, since macroscopic light consisting of many photons, behaves like a wave to an excellent approximation, whereas individual photons make tiny spots on photographic film and clicks in detectors, behavior more consistent with the particle picture. However, when interpreting the photon as point particle, we run into irreconcilable logical difficulties when interpreting single- and double-slit experiments with photons. We are therefore forced to assign both wave and particle aspects to photons. (Even electrons, the best examples of point-particles in nature, sometimes behave as waves to such an extent that even scientific instruments, such as the electron microscope, would not work, were it not for wave-particle duality). According to this principle: A quantum object (a) is produced as a **particle**, (b) propagates like a **wave**, and (c) is detected as a **particle** with a probability distribution that corresponds to a **wave**.

The Machinery of Quantum Mechanics

There are two layers to quantum mechanics: The conceptual and the machinery. We review many baffling quantum effects and learn how to interpret them with the help of the three fundamental rules of quantum mechanics: Feynman's Rule, Born's Rule and the Composition Rule. While these three rules are sufficient to explain many quantum effects qualitatively, we need the machinery of quantum mechanics to make quantitative predictions. The most powerful tool in the quantum shed is Schroedinger's equation. Given a quantum system, Schroedinger's theory allows us to formulate an equation for the amplitude of a system. Solving Schroedinger's equation, we obtain this amplitude explicitly, and with it possess the maximal information about the system allowed by quantum mechanics. Since all the information about the system is encoded in its wave function, the natural question to ask is how to extract this information and how to predict the values of physical observables such as position and momentum on the basis of the wave function. This question can be answered when we connect the formal, mathematical machinery of quantum mechanics to the real world. This is accomplished by assigning Hermitian operators to physical observables. The operators may then be used to "operate" on the wave function to extract physical system information. We then study several examples of how to do this in practice. We also introduce and study the important concept of the spectrum of a quantum system. We learn that the spectrum is computed by solving certain eigenvalue equations that directly derive from Schroedinger's equation. The model systems studied next are carefully

selected to exemplify the most important type of spectra encountered in real-life quantum systems. Next we take a closer look at Dirac's notation already introduced in the previous chapter. Not only do we develop it into a handy tool to perform actual quantum calculations, we also encounter many parallels to ordinary linear algebra, which makes Dirac's notation more palatable, since it suggests that Dirac's notation is nothing but a convenient adaptation of linear algebra to the specific needs of quantum mechanics. We then proceed to make extensive use of Dirac's notation when we discuss Heisenberg's formulation of quantum mechanics. We also formulate some new operator-based techniques for the solution of Schröedinger's equation. The most useful one is the time evolution operator $\widehat{U}(t)$, which also plays a central role in the transition from Schroedinger's formulation to Heisenberg's formulation of quantum mechanics. Next, in order to prepare the ground for quantum computing, we study quantum two-level systems. The central point here is to show that we have complete control over two level systems. This means that if we have a physical means to induce transitions from the ground state to the excited state of a two-level system, for instance a magnetic field or a laser, then we are able to produce any two-level quantum state whatsoever, at will. This is a fundamentally enabling technology for the construction of quantum computers.

Measurement

In classical physics the theory of measurement plays a minor role and is not usually emphasized in textbooks on classical mechanics. The reason is that in classical physics it is assumed that measurement is a process that can be minimized to such an extent that it negligibly influences the dynamical state of the system. In contrast to the subordinate role of measurement in classical physics, measurement plays a central role in quantum mechanics whose influence on the system to be measured is usually large. Moreover, the act of measurement itself is poorly understood and to this day measurement is the most controversial part of quantum mechanics. Sidestepping the controversy, we adopt von Neumann's view of measurement as a nondeterministic collapse of the wave function, a process strictly outside the realm of Schrödinger's equation that requires its own axiom of measurement. We argue that even if someday the act of measurement is explained microscopically as a purely deterministic effect amenable to a dynamical description via the deterministic Schrödinger equation, this would not diminish the value of the von Neumann picture of measurement. The relation between the "new," microscopic theory of measurement and the "old" von Neumann theory would be akin to the relationship between statistical mechanics and thermodynamics, where statistical mechanics plays the role of the microscopic theory that explains the effective, phenomenological theory of thermodynamics without invalidating it. We next turn to one of the most profound differences between classical and quantum mechanics. While in classical mechanics all system variables can *always* be measured simultaneously with arbitrary accuracy, this is not the case in quantum mechanics. Heisenberg's Uncertainty Principle imposes a fundamental limit on the accuracy with which the values of two incompatible observables can be measured simultaneously. Using our quantum machinery developed in the chapter before, we're able to derive the exact formulation of Heisenberg's Uncertainty Principle without making any assumptions. We then encounter another counterintuitive quantum effect: It is impossible to make copies of an unknown quantum state. This is known as the quantum no-cloning theorem and, again, we can prove it in full generality. We next study yet another quantum effect intimately connected with quantum measurement: The quantum Zeno effect. Here, a sequence of measurements is used to slow down the time evolution of a quantum system. The quantum Zeno effect is one of the most fascinating quantum effects. It has already found a practical application in connection with interaction-free measurements. We will see it again soon.

Interaction-Free Measurements

Sounds like a contradiction in terms and yet it works. A first hint of how to perform interaction-free measurements was provided by Renninger in the 1950s. He noticed that the absence of detection is

sometimes a measurement, too, which may provide valuable information on location and momentum of a particle. Renninger called this type of measurement negative-result measurements. We illustrate this idea in connection with the familiar double-slit experiment. While this setup provides the conceptual foundation for interaction-free measurements, it is not of much use in practice. A first practical scheme for interaction-free measurement was provided by Elitzur and Vaidman who proposed using a Mach-Zehnder interferometer to perform interaction-free measurements. We study their scheme in detail. It illustrates that interaction-free measurement plays on the wave-particle duality. The wave aspect of light allows us to zero out the signal in one of two detectors. The particle aspect of light forces photons to take one of two paths in case an observer is present in one of the interferometer arms and destroys wave coherence. The ideas of Elitzur and Vaidman were implemented experimentally by Kwiat and collaborators with the help of a Michelson interferometer. We study their experimental scheme. Once the idea of interaction-free measurements had been confirmed experimentally, Kwiat and collaborators went further. With the help of an ingenious optical implementation of the quantum Zeno effect they obtained interaction-free observations exceeding 50% efficiency. Meanwhile interaction-free measurements evolved into interaction-free imaging of small objects. Interaction-free imaging may find useful applications whenever an object to be imagined is particularly sensitive to the light.

The EPR Paradox

Although Einstein helped create quantum mechanics in the years before 1925, he never accepted the "final product," the quantum mechanics of Heisenberg and Schrödinger, created in 1925-1926. Einstein was particularly displeased with Born's probabilistic interpretation of quantum mechanics, and Heisenberg's Uncertainty Principle, which both imply that quantum theory, unlike classical mechanics, is unable to predict simultaneously exact values of the positions and momenta of particles. In addition, Einstein thought it untenable that reality should be tied to measurements, and that, according to quantum mechanics, a local, objective reality may not even exist between quantum events and measurements. Einstein's dissatisfaction with these "deficiencies" of quantum theory is only natural. Einstein was a classical physicist at heart. He lived in the Newtonian tradition where the attributes of particles, such as positions and momenta, are always objectively real, i.e., existing independently of observers and measuring devices. In fact, with his special and general theories of relativity, he completed the classical worldview. Therefore, fully accepting the successes of quantum mechanics within its range or applicability, he was dissatisfied by what he considered to be essential limitations of quantum mechanics, namely its probabilistic interpretation, and the inability of quantum mechanics to predict sharp results for physical observables that correspond to noncommuting quantum operators. As a consequence, Einstein thought that quantum mechanics is an incomplete theory that is badly in need of improvement. In order to understand what EPR mean by an "incomplete theory," we study theory building and the properties of physical theories. Einstein's best stab at proving the incompleteness of quantum mechanics occurred in 1935 when he published a paper co-authored with Podolsky and Rosen, the "EPR paper". In it, EPR present an ingeniously constructed scattering system that seemingly leads to the paradox of being able to predict sharp values of noncommuting physical observables. This is referred to as the "EPR paradox". We review this part of the EPR paper. However, analyzing in detail the collapse of the EPR wave function, we conclude that EPR's argument is flawed. EPR assigned a definite quantum mechanical state to subsystems of the EPR system that are entangled, and therefore not in a definite state at all. Understanding this part of quantum mechanics resolves the EPR paradox, and leads to a better understanding of quantum systems, subsystems and entanglement. Thus, although ultimately flawed, the EPR argument contributed decisively to sharpening the concepts and the language of quantum mechanics. Next we present Bell's analysis of the EPR paradox. Reaching far beyond qualitative, philosophical arguments, Bell proved quantitatively, using mathematical derivations that the local realism that

EPR favored, simply does not exist as a universally valid concept in nature. Bell showed this by proving that local hidden-variable theories, possible implementations of Einstein's "local reality program," are incorrect descriptions of nature. More than an elaboration of an obscure, technical point in quantum mechanics, by proving that local reality is not a fundamental feature of nature, Bells work is as revolutionary as Einstein's relativity and quantum mechanics themselves, and marks one of the greatest scientific revolutions of 20th century physics. An illustration of the EPR "paradox" is Mermin's Reality Machine, which we present. It allows a more intuitive and more direct appreciation of the problem of local reality in nature.

Classical and Quantum Information

Now we lay the conceptual foundations of quantum computing and quantum information processing. Although the computational power of classical computers is impressive, increasing their performance relies on continued miniaturization of classical electronic circuitry. First we see that due to the atomistic nature of matter this miniaturization cannot go on forever, setting distinct fundamental limits to classical computing. Therefore, a further increase in computing and information processing power can be achieved only if we change the computing paradigm. Quantum computing and information processing provides an example of such a paradigm shift. It promises unprecedented computer power, computer power so tremendous that it easily exceeds the computer power of a classical computer the size of the universe! While classical computers are based on binary switches, i.e., bits, that can take only the logical values 0 and 1, quantum computers are based on qubits that may be in a quantum superposition of $|0\rangle$ and $|1\rangle$, which accounts for one of the "secrets" of quantum computing. We take a look at classical bits and quantum qubits. In order to perform a computation, no matter whether classical or quantum mechanical, we need gates that allow us to perform logic operations on bits (classical gates) or qubits (quantum gates). We study classical and quantum gates together with their main differences. Next we combine classical and quantum logic gates into classical and quantum circuits. Now we encounter our first meaningful quantum circuit, the quantum register loading circuit. This circuit is a common element in quantum algorithms. In fact, without it, it is hard to imagine how we could possibly make effective use of quantum parallel processing and obtain the exponential speedups that quantum algorithms afford us. Next we learn how to teleport the quantum state of a single electron. Although teleportation of the quantum state of a macroscopic object is impossible to do with present-day technology, there are no physical laws that would forbid it. While teleportation has obvious technical applications that even include the possibility of a quantum internet (discussed in the last chapter), its main scientific value is of a conceptual nature: it shows how to resolve the quantum information of a given single-qubit state $|\phi_1\rangle$ into two components, two bits of classical information and a purely quantum EPR correlation.

Quantum Computing

This is where we start studying the basic ideas behind quantum computing. We start by constructing a primitive quantum computer that illustrates how superposition may be used to achieve parallel processing in quantum computers. We also see that parallel processing is not enough: It has to be combined with interference in order to select desired results from the multitude of computed results. Deutsch's algorithm was the first to combine both principles. Deutsch's algorithm performs a task that cannot be performed on any classical computer in principle. Thus, Deutsch's algorithm was the first to prove the point that quantum computing is a qualitatively new way of information processing. While Deutsch's algorithm beats any classical algorithm by a factor of two, practically speaking, this is not too impressive given how difficult it is to construct quantum circuitry and keep its coherence. Simply using two off-the-shelf classical processors and running them in parallel, erases the speed advantage of Deutsch's algorithm. This is where the Deutsch-Jozsa algorithm comes in: It demonstrates that quantum computers can perform computations that are out of the league of classical computers, even if we allow for a classical computer the size of the universe! Therefore,

when it comes to solving problems of the Deutsch-Jozsa type, classical computers simply are no match for quantum computers. Although quantum computers excel when solving problems of the Deutsch- or Deutsch-Jozsa type, one could argue that these problems and their quantum solution algorithms are rather contrived and do not have any practical applications. To counter this argument, and to show that quantum computers reign supreme even in areas of everyday importance, we study Grover's quantum algorithm. Grover's algorithm addresses the problem of finding an item in an unsorted database. Not only since the advent of the Internet and its various search engines do we know that searching for items in an unsorted environment is an everyday occurrence. We see that Grover's algorithm solves the task of locating an object in an unsorted database of size N^2 in approximately N steps, while any classical algorithm requires $N^2/2$ steps, on average. This may seem like a marginal advantage but as the size of the database (haystack) increases the savings are more and more significant. Based on the examples of quantum algorithms presented in this chapter, it is not hard to imagine that quantum computers have the potential to revolutionize the fields of information processing and computing. In the next chapter we examine their potential to disrupt.

Classical Cryptology

Secure communication is an enabling technology for all of modern military, government, and commercial information exchange. Internet commerce, and especially Internet banking, are unthinkable without convenient cryptosystems for the secure transmission of business transactions. There are essentially two different ways to set up a cryptosystem for the secure exchange of messages: symmetric, private-key cryptosystems, and asymmetric, public-key cryptosystems. While private-key cryptosystems display the highest degree of security, they are awkward and inconvenient in practice and suffer from the key-distribution problem. Public-key cryptosystems, in particular the RSA cryptosystem, are the dominant technology in today's Internet applications. RSA, however, is only as safe as integer factorization is difficult to do. Advances in number theory, resulting in a fast factoring algorithm, may render RSA obsolete. But even in the absence of number-theoretic breakthroughs quantum computers have the potential to crack RSA. A quantum algorithm specifically developed for cracking RSA is Shor's algorithm which we discuss next.

Quantum Factoring

According to present scientific consensus, classical computers cannot crack RSA cryptosystems with RSA moduli that have 300 decimal digits or more. The reason is that classical factoring algorithms requires an execution time that grows exponentially with the number of digits of the RSA modulus. Quantum computers, however, have the capability of executing exponentially many instructions in parallel, thus outclassing any classical computer. However, a quantum computer without quantum software is powerless. In this chapter we present the basis for a possible future quantum factorization software, Shor's algorithm, that has the ability/potential to crack the RSA cryptosystem. It is based on an ingenious combination of the "classical" Miller algorithm for factoring semi-primes and the quantum Fourier transform (we discuss both).

Ion-Trap Quantum Computers

In theory quantum computers are powerful devices. But do they work in practice? In this chapter we show that the answer to this question is an emphatic: "yes"! Nothing is more convincing than demonstrating a working quantum computer in the lab. And several experimental groups throughout the world have done just that (this is 2009): These groups are operating quantum computers that are based on various quantum computer architectures, one of them the ion-trap quantum computer design. Due to its conceptual simplicity we chose the $^{40}Ca^{+}$ ion-trap quantum computer of the Innsbruck group as a representative example. The centerpiece of the Innsbruck quantum computer is a linear radio-frequency ion trap. Initially, the trap is loaded with several $^{40}Ca^{+}$ ions, which are dynamically confined to the trap, but form a relatively "hot" (i.e., room temperature) charged gas, a non-neutral plasma. In this state the ions are useless for quantum computing. However, application of a cooling laser, reduces the kinetic energy of the ions, eventually lining them up on

the axis of the trap. This state of the ions is sometimes referred to as the crystalline state of the $^{40}Ca^{+}$ ions, in which each $^{40}Ca^{+}$ ion represents one qubit of the quantum computer. As far as the hardware is concerned (the trap and the ions), the qubits are now ready for a quantum calculation. This is where the microprogramming, i.e., the realization of basic quantum gates, comes into play. According to the Cirac-Zoller scheme, laser pulses and the ion-crystal's center of mass motion, are used to implement the quantum CNOT gate. Since the CNOT gate is universal (with Hadamard), all quantum computations can be done on a quantum computer once the CNOT gate is implemented. Thus, the trap hardware together with the Cirac-Zoller scheme, implements a universal quantum computer. The details of how this is done in practice are presented here, providing positive proof that quantum computers are now a reality[1].

Outlook (Commencement)

Not long ago it was inconceivable how pure research in the foundations of quantum mechanics could possibly lead to practical applications. Yet, here we are, with Einstein's "spooky action-at-a-distance" providing the foundation for a completely new industry, quantum information technology. The relatively short time it took to turn quantum paradoxes into practical quantum devices provides one of the success stories of how curiosity-driven research results in a major scientific and industrial revolution.

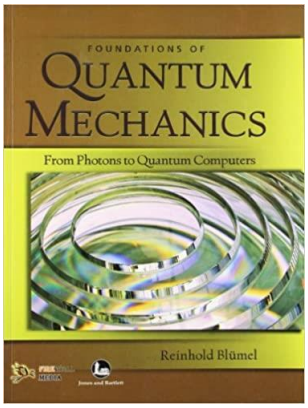

[1] See, e.g., hardware like in the Richerme lab. Other modalities: **topological qubits** (belong to the pure, or exotic, superconducting modality) and seem to be our only chance to get out of the NISQ range. Current NISQ modalities: **superconducting qubits** (fast, technology (silicon wafers, CMOS) is well established, but integration of control and readout technologies that maintain qubit coherence even at mK temperatures remains a challenge). **Trapped ions** are very stable systems, and very well characterized (technology has been in use in atomic clocks for decades) but main challenge is 3D integration. **Doped Si qubits** leverage silicon fabrication technology and spin qubits have long coherence times but they're rather slow and scalability/integration remains an issue. **NV centers** afford both an electron spin and a nuclear spin, and have large(r) coherence times (with lower coherence they can even be operated at room temperature) but scalability remains an issue (hard to place in precise locations to create large qubit arrays). **SiGe quantum dots** leverage well established technology, are relatively small and controlled by gate voltages; their main challenge, again, is needed 3D integration techniques. **Neutral atoms** make highly stable qubits with coherence times that are comparatively very long; this modality will require integrated optics to ultimately scale up. Some of the other modalities that could be mentioned here: **linear optics** (hard to make photons interact with one another), **molecular ions**, **electrons on liquid helium** and (again) **Majorana fermions** (top choice since they will exhibit, crucially, topologic protection to errors).

Final thoughts (for now): In 2006 Vazirani published a book with Dasgupta and Papadimitriou. It starts with an ancient hard problem (factoring) and ends with a chapter on quantum computing (and its promise to disrupt). The book asks (on p. 310, this is 2006): "Can quantum computers be built?" Ten years later IBM Q was launched, as was Rigetti. Online courses (professional certificates in quantum engineering) started to appear too at about the same time. If one considers how pervasive a sense of astonishment in fact is (e.g., David Deutsch himself in [2]: "Last year I saw their ion-trap experiment, where they were experimenting on a single calcium atom […t]he idea of not just accessing but manipulating it, in incredibly subtle ways, is something I totally assumed would never happen. Now they do it routinely") one immediately understands why the probability remains significant to continue being pleasantly surprised by the pace of advancements in technology. For the most skeptical though it's worth sharing the brilliant argument of Michele Mosca [1] that addresses the time frame aspect from a somewhat inverted perspective: if we add the security shelf life (of our secure data) to migration time and then compare to the collapse time any responsible manager or executive immediately feels the impending doom (or, rather, sense of urgency).

[1] https://www.youtube.com/watch?v=vipU_-QGoOg (Michele Mosca relevant part @28:32)
[2] https://www.economist.com/technology-quarterly/2017/03/09/david-deutsch-father-of-quantum-computing

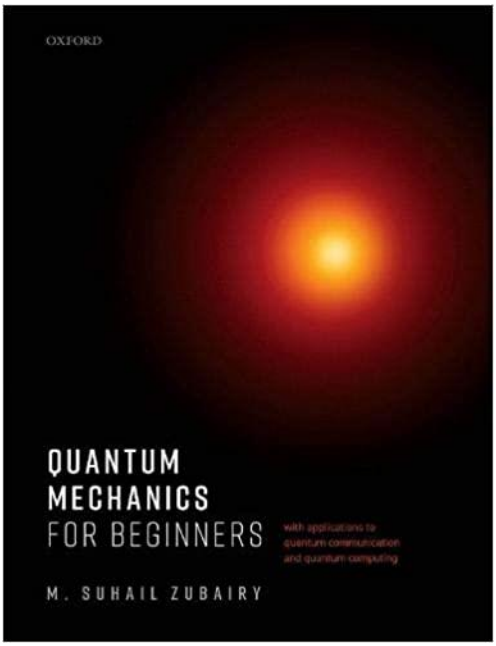

Possibly alternate reference (published July 2020):

M Suhail Zubary, Quantum Mechanics for Beginners (with applications to quantum communication and quantum computing), published by Oxford Univ Press, July 2020.