The Great Concepts:
A Focus on Creation and Knowledge

by

Professor Sima Dimitrijev
School of Microelectronic Engineering

Professorial Lecture

Delivered at

GRIFFITH UNIVERSITY
Nathan Campus

18 November 2004

Professorial Lecture

by

Sima Dimitrijev
School of Microelectronic Engineering
Griffith University, Nathan, Qld. 4111

Prologue

A professorial lecture is expected to present research achievements. My formal and published research activities have been focused lately on the potential to develop the next generation semiconductor memories — nonvolatile dynamic memories on silicon carbide. It is a fascinating topic, especially because of its potential to make an enormous impact on information systems. I could use analogies and metaphors to present in this lecture the main characteristics of a semiconductor memory and then to explain our memory approach. As convincing as this presentation might be, the main discussion after that would inevitably focus on the future: what should be done by the technical and business management of the project, what decisions may be taken by potential investors, what actions may be taken by competitors (some of them real gorillas on the market),... With a strong focus on future prediction, it would be very difficult for me to convince anyone in a sentence that it can be far easier to create the future than to predict it.

Instead of semiconductor metaphors, I decided to set this lecture at the more general level of metaphors themselves and related concepts. I decided to share with you my more general way of thinking that, for example, leads me to the conclusion that it can be easier to create the future than to predict it. Being unpublished material, this presentation should not be regarded as a formal and rigorously performed research. It has never been prepared for a peer review. Nonetheless, this does not diminish the importance of the presentation. The topic of this presentation is as important as any decision we make. An example may be my decision to not only start working but also to persist with the research in silicon-carbide electronics. Being general, however, the logic presented in this lecture does relate to any decision we make. Furthermore, the complete operation of our conscious brain, including the role of knowledge, can be seen as a decision-making process. As will be shown, the decision making is a sophisticated mechanism for a selection of available options: a consciousless selection is referred to as the natural selection, whereas a knowledge-based selection is the decision making. Accordingly, there is an accepted and a rigorous scientific foundation for the thinking that will be presented — the theory of evolution.

The Great Concepts

A Focus on Creation and Knowledge

Sima Dimitrijev

Contents

1 Introduction

This lecture is about a search for fundamental *ideas* that form the basis of reasoning and action. In the existing scientific paradigm, the correct way of saying this would be that the lecture is about fundamental *principles*, given that a typical definition of *principle* is “a fundamental truth or law as the basis of reasoning or action”. Yet, reasoning and actions of people do not blindly follow strict principles and can be unpredictable. There is a difference between fundamental ideas, that is *concepts*, and absolute truths or laws. The thinking presented in this lecture is based on somewhat flexible *concepts* rather than on absolute *principles.*

There are both principles and concepts that do not work outside a strictly defined framework. Concepts that cannot be taken out of a physics textbook and used to explain biological phenomena are not fundamental for both areas. These are not great concepts. There are no natural boundaries between physical, chemical, biological, psychological, and even sociological phenomena. This fact motivates the search for concepts that are general enough to have unique meaning in any relevant area of knowledge. These interdisciplinary and interlevel concepts are the *great concepts.*

Levels here mean *degrees of complexity.* The bottom level is associated with physical phenomena, including both small-size effects described by quantum mechanics and large-size objects that can be approximated by Newton’s mechanics. At a higher level of complexity, we find qualitatively different phenomena, such as reproduction of plants and animals. At even higher level of complexity, we deal with the phenomenon of consciousness. Although the concepts from a high level of complexity are not applicable to simple systems, the great concepts applicable at low levels maintain their meaning and even their importance at high levels of complexity. This postulate means that the effects of a great concept cannot be limited to the mysterious quantum-mechanical level but there should be related effects that we observe and deal with in our everyday lives. It means that we can use an approach that is analogous to the very successful use of metaphors and analogies as communication and teaching tools.

2 Order

Order is a great concept because it appears prominently at all levels — from the subatomic phenomena to human abstractions. At the level of human abstractions, a good example for order are the Maxwell’s equations of electromagnetism. These are deterministic differential equations that model the space and time variations of electric and magnetic fields. Maxwell's equations are in every textbook on electromagnetism, which is an illustration of a high degree of order as far as the theory of electromagnetism is concerned.

At a lower level, the world of animals presents profound examples of order. The examples are not limited to the animals themselves, but extend to their creations, such as honeycombs and spider webs. The order associated with plants, and in particular with flowers, is not only apparent but admirable.

Order is not limited to living organisms. The example of crystals shows that order is observed even at the atomic level. Furthermore, the periodic table of elements shows that there is order at the subatomic level as well.

2.1 Replication

Replication is a simple mechanism that creates order. At the atomic level, crystals are orderly structures that are created by replicating the pattern of a unit cell. Many crystals are created by humans, but there are many that were created well before the emergence of humans. In fact, it is possible that the evolution of humans originated from simple crystal structures. A description of this quite appealing possibility appears in *The Selfish Gene* by Dawkins[1], although it is not his preferred theory. Replication is the central topic of this book that provides powerful descriptions of the importance of replication mechanisms. It demonstrates far reaching consequences of a simple fact: *replication leads to prominence of what is being replicated.*

In the existing scientific paradigm, we can relate the mechanism of creating order by replication to the *minimum-energy principle.* Under certain conditions, a specific pattern of atom placement provides the minimum-energy situation, so it is this pattern that is replicated to spontaneously create a crystal consisting of billion billions of atoms per mm^3 of material. Humans would not be able to place that many atoms into a crystal pattern without the spontaneous replicating mechanism.

The minimum-energy idea can be extended to the level of human knowledge. The complex systems designed by humans — such as computers, airplanes, buildings,... — are never designed from scratch. To minimize the energy for a given design, or to maximize design performance with the available energy, every human design is based on replicated knowledge that is either adapted or slightly upgraded. This description also applies to the consciousless designs of nature, such as the existing plants and animals that evolved through natural selection of slightly modified replicas of successful genes.

2.2 Generalisation

Generalisation is a mechanism that creates a higher level of order than replication. Generalisation is about identifying common characteristics in otherwise nonidentical objects and phenomena. It can metaphorically be described as distilling principles or seeing common patterns and relationships in an apparent chaos. Equivalently, generalisation is about ignoring irrelevant specifics. This description of generalisation is important because of its focus on the *irrelevant* specifics that are *ignored.* Irrelevant specifics for some conditions or purpose may appear as central elements in different conditions or for different considerations. This simply means that no principle derived by generalisation can be absolute.

An example for order created by generalisation is the Newton's second law of motion: $F = m(dv/dt)$ where F is the force applied to an object of mass m and dv/dt is the acceleration of the object (change of velocity dv per unit time dt). Objects may have different dimensions, different shapes, different colors, can be made of different materials, but all these are irrelevant specifics. The only characteristic of an object that matters for the relationship between force and acceleration is the mass of the object. This law creates a high-degree of order in terms of motion of as different objects as air molecules, billiard balls, and planets. But it cannot be an absolute law because it does ignore specifics such as the dimensions and shape of an object (this was, of course, demonstrated by the development of quantum mechanics — the theory that explains the behaviour of small things).

The specific characteristics that are ignored in the process of generalisation set the limits to both the applicability and the meaning of the generalized law. The Gibbs paradox illustrates this. This paradox relates to the calculation and the interpretation of *entropy*, the central concept of the second law of thermodynamics, but also the central concept in the theory of information. For an isolated thermodynamical system in equilibrium, the entropy is at its maximum and can be related to the heat normalized by the average thermal energy of the individual particles. For systems that are not in equilibrium, the entropy is a measure of the degree of *order* in the reverse sense: a higher entropy corresponds to a lower degree of order. The second law of thermodynamics then states that spontaneous thermal processes increase the entropy as they work to bring the system into thermal equilibrium. In other words, the spontaneous thermal processes diminish the degree of order. This enables to present the following adaptation of the Gibbs paradox.

> Imagine randomly moving particles (small circles) in a box as a model of a gas in a closed room. Assume the box can be partitioned

> in two parts by an inpenetrating wall. For a system in equilibrium, we do not observe any change in the degree of order (disorder) as the wall is inserted and removed. The entropy does not change. Assume now that the wall is inserted and that the colour of the particles in one part is changed from green to red. Next, remove the wall to allow the random motion to mix the red and green particles. The initial degree of order (the red particles in one part and the green particles in the other part of the box) is gradually diminished until they are fully mixed. The entropy has increased to a new maximum value at equilibrium. This mixing-gas paradox is due to the difference in maximum entropy for the system with a single type of particles (a single colour) and particles with different colours.

What does the colour of particles have to do with the maximum entropy, that is, with the heat normalized by the average thermal energy of the individual particles? Clearly, the answer is — nothing. When considering order/disorder within the limits of thermal processes, the colour should be ignored as an irrelevant specific. When calculating the entropy within thermodynamic limits, the particles should be considered indistinguishable regardless of the fact that it may be possible to characterise them with different colours or labels such as "oxygen" and "nitrogen". Otherwise, the meaning of diminished order — and the associated entropy increase — is taken beyond the generalisation limits of the second law of thermodynamics with the illusion that the colour of the particles has something to do with thermal processes. In analogy with this, many philosophical conclusions derived from the second law of thermodynamics are also just illusions.

Another example of indistinguishable objects is males and females as programmers. The idea that a good or bad performance has nothing to do with the sex of the programmer is a perfectly valid generalisation (a removal of irrelevant specifics). Yet, by defining and focusing on *specific* programming-related skills or characteristics, it is possible to establish legitimate correlation to the sex of programmers. No generalized law is absolute.

2.3 Determinism and Deterministic Equations as Approximations of Reality

Since Newton, the favorite way of expressing general laws in physics is in the form of differential equations. The second Newton's law is an example of a simple differential equation: $F = m(dv/dt)$. Differential equations provide another example for generalisation that can be taken too far — to the level of abstraction that is just an illusion.

Take as an example even simpler differential equation than the Newton's second law, a differential equation that relates two fundamental concepts of physical reality: energy and time. A simple example for energy is the electrical energy consumed by a household. As the consumed energy continuously increases with time, it becomes more convenient to introduce the concept of energy consumption per day or per hour, or per time interval dt in general terms. Thus, energy consumption dE per unit time dt is defined as power consumption: $P = dE/dt$. If the time interval dt is too long, the energy consumption may fluctuate within the time interval. For example, the energy consumption of a household changes during the day. If dt is taken to be one day, then $P = dE/dt$ does not seem to represent the real power but just an average power — a kind of approximation that does not seem precise enough. To express the real power, dt has to be small enough so that there is no fluctuation of energy consumption during dt. In other words, dt has to be small enough so that the power consumption remains constant during dt. A time interval that satisfies this condition is referred to as infinitesimal. Now, consider the case when the power consumption is suddenly reduced to zero because the main switch is turned off. Unless the power consumption decays in a staircase fashion with known step size(s), there are no dt values other than zero that will ensure that P is constant during dt. However, the case of $dt = 0$ does not provide a realistic model for any change of energy in time. Therefore, dE/dt can only be an approximation and not an absolutely precise and determined variable.

Any differential equation with time as a variable is deterministic. This is because its solution, set for specific initial conditions, predicts the future values of the equation's variables without any room for uncertainty. Consequently, an acceptance of a differential equation as absolutely correct model of reality is inseparable from the acceptance of absolute determinism[1]. The alternative to this scenario is to consider the models based on differential equations as approximations of reality rather than the absolutely correct models. Of course, approximations can be very good, but then, they can be very poor in some circumstances.

---

[1] One exception is the Schrödinger's wave equation, which is a deterministic differential equation, but its variable itself involves some uncertainty.

3 Uncertainty

There is no perfect order: crystals always have defects, flowers are not perfect, and the deterministic equations do not predict the future with absolute certainty. Although future moon eclipses can be predicted with high accuracy, there is no certainty these predictions will hold forever. An unpredictable collision between the moon and a large meteorite cannot be ruled out.

The concept of order that is not perfect is not too hard to accept. The ultimate reason for the lack of perfection, however, has presented unsurmountable challenge to even the greatest thinkers. A frequent answer is that there are fluctuations in either the initial conditions or the parameters of otherwise deterministic processes. The question then is what causes the fluctuations. Not many people would seriously consider the idea that some fluctuations can occur without a direct cause. Nonetheless, it is a fact that the alternative — the absolute causality — is inseparable from absolute determinism, which negates both the idea of fluctuations and the concept of imperfect order.

3.1 The Limit of Causality (Schrödinger's Cat)

Erwin Schrödinger's famous wave equation became the widely accepted mathematical model for the behaviour of small things, which is the subject of quantum mechanics. Schrödinger's equation is a deterministic differential equation, but its variable — the wave function, $\psi(x, y, z, t)$ — involves uncertainty. It is generally accepted that $\psi\psi^*$, where ψ^* is the complex conjugate of ψ, is the probability of finding the object described by its wave function at point (x, y, z) at time t. It is the meaning of the wave function that has caused real problems for the developers of quantum mechanics, including Schrödinger himself who expressed his feelings about quantum mechanics by the following sentence: "I do not like it, and I'm sorry I ever had anything to do with it". In an attempt to demonstrate that something has to be wrong with the interpretation of the wave function that is now known as the Copenhagen interpretation, Schrödinger proposed his famous cat paradox. A translation of Schrödinger's description is as follows:

> ... For example, consider a cat is penned up in a steel chamber, along with the following diabolical device (which must be secured against direct interference by the cat). In the device is a Geiger counter with a tiny bit of radioactive substance, so small that perhaps in the course of one hour only one of the atoms decays, but also, with equal probability, perhaps none. If the decay happens, the counter tube discharges and through a relay releases a hammer

> that shatters a small flask of hydrocyanic acid. If one has left this entire system to itself for an hour, one would say that the cat still lives if meanwhile no atom has decayed. The first atomic decay would have poisoned it. ...

Before considering the paradox related to the Copenhagen interpretation, it is helpful to focus on Schrödinger's cause-and-event chain. Although somewhat long, the cause-and-event chain in this example is straightforward: if the cat is found dead, the cause of the death is the poison spilled from the broken flask; the cause for the broken flask is a hit by the hammer and the action of the hammer is caused by a relay movement; the relay movement is caused by a signal from the counter tube, and the signal in the counter tube is caused by a radioactive decay. The big question is about the cause of the radioactive decay. It may seem natural to assume that an imminent cause for the radioactive decay must exist in principle, regardless of the fact that it has not been discovered. Importantly, this assumption just shifts the question of the *initial* cause one step further — to the cause of the event that triggered the radioactive decay. Applying the assumption that the possible cause for the radioactive decay is caused by something else and that "something else" is also caused by something else leads to an infinite cause-and-event chain. Clearly, there can be neither initial event nor initial cause if the chain is to continue to infinity.

The concept of *infinite* cause-and-event chains removes, at least in theory, the need to deal with the really difficult question related to the *initial* cause. For some thinkers, it may indeed be easier to accept the concept of infinity than the concept of initial cause. The problem for these thinkers is, however, that the cause-and-event chain hits its end at the quantum-mechanical level if $\psi\psi^*$ is genuine probability. Many attempts have been made to find a meaning of the probability $\psi\psi^*$ that avoids the concept of *random initial events.* The most accepted attempt is the Copenhagen interpretation of quantum mechanics. It is this interpretation that becomes paradoxical when applied to Schrödinger's cat.

In a nut shell, the relevant aspect of the Copenhagen interpretation is that the wave function of a system is a superposition of all the possible options, a specific option being selected upon *observation* of the system — so called wave function collapse that is *caused* by an observer. According to the Copenhagen interpretation, the wave function of the system with the cat would have to involve both the living and the dead cat smeared out in equal parts. The problem is that the wave-function collapse (the selection of the *dead* or *alive* state of the cat) is caused by an observer who opens the chamber to inspect the cat, yet it could easily be established that the cat died perhaps more than half-an hour before the chamber was opened by the observer.

The idea that an observer is needed to *cause* the wave-function collapse is a generalisation, so it can legitimately ignore *irrelevant* specifics. The cat paradox does link quantum phenomena (radioactive decay) to a complex biological system, but in doing so, it does not bring the biological system as an irrelevant specific system that can be ignored. If the only cause that can be found for the collapse of the wave function that corresponds to the radioactive decay is an interference by an observer, then this is the only possible cause for the death of the cat. Yet the fatal radioactive decay could have occurred well before the interference by the observer. Not surprisingly, the idea about the essential role of an observer has been confronted with many logical problems (Gribbin [2] provides a brief review of these problems).

There is no paradox, however, with the simple postulate that a radioactive decay is a random event. Neither an observer nor any other cause is necessary for the collapse of the relevant wave function. So, if an event of radioactive decay killed the cat half-an-hour before the observer opened the chamber, then a valid post-mortem examination will be in agreement with this timing and not with the time of chamber opening.

The concept of genuine random events has been very hard to accept. In addition to the Copenhagen interpretation, which remains the most accepted attempt to avoid this concept, an illustrative example is Einstein's widely cited rejection of this idea: "At any rate, I am convinced that He [God] does not play dice". In spite of the respectful disbelief of many big thinkers, it is a simple fact that the idea of random events does not lead to contradictory conclusions. Moreover, it appears as a great concept with abundance of related effects on display: from the quantum-mechanical level (such as the radioactive decay), through the impressive world of plants and animals *created* by natural selection of random mutations, to human creativity.

3.2 Uncertainty Principle in Quantum Mechanics

The concept of genuine uncertainty, or genuine randomness, is not to be extrapolated to a perfect disorder. Genuine uncertainty does not mean total uncertainty in the sense that every change is totally unexpected. The concept is that there is some degree of genuine uncertainty, in full agreement with the concept of imperfect order.

The idea of a *partial* genuine uncertainty may seem paradoxical. It rises the question of the elementary unit of uncertainty — the smallest possible uncertainty before it disappears altogether. This issue is the same as in the case of the *infinitesimal*. Both the *unit of uncertainty* and the *infinitesimal* should not be so small that they do not exist, but how small can they be? Any answer to this question appears absurd if we think in terms of continuous sizes, as in the case of differential equations. If a size reduction is to be associated with the infinite chain of real numbers decaying toward zero, then there is no natural unit. However, there is no need to think in terms of continuous sizes. We can adopt the model of quantized sizes. If we do this, we are on the territory of quantum mechanics (Richard Feynman provides an insightful introduction to quantum mechanics [3]). The unit of uncertainty can be related to the Heisenberg's *uncertainty principle*, which is the fundamental principle in the presentation of quantum mechanics that is an equivalent alternative to Schrödinger's wave-function approach.

The uncertainty principle links together two physical quantities whose product is in the units of *action* — Js. The obvious example is *energy* (J) times *time* (s). Applied to these quantities, the uncertainty principle can be expressed by the following *inequation*:

$$\Delta E \Delta t \geq h$$

where $h = 6.625 \times 10^{-34}$ Js is the Planck constant, ΔE is the uncertainty in energy and Δt is the uncertainty in time. Another pair of physical quantities whose product is in the units of Js is *distance* (x) and *momentum* (p_x):

$$\Delta x \Delta p_x \geq h$$

The energy–time inequation can be applied to *processes*. An example is when the electrons in hydrogen atoms are excited to the first higher energy level by colliding the hydrogen atoms with a beam of electrons. The hydrogen electrons remain on the excited level for an uncertain time period Δt. This uncertainty relates to uncertainty (fluctuation) ΔE in the energies taken from the initial monoenergetic electron beam. Another example is the radioactive decay in the cat paradox.

The distance–momentum inequation can be applied to *objects*. A good example is an electron whose position is not specified with a better certainty than Δx if the momentum uncertainty is Δp_x. With many electrons, the probability distribution for their positions can appear as the intensity of an electron wave. An example of electrons acting as waves is the scanning electron microscope.

Based on Heisenberg uncertainty relations, it can be assumed that the Planck constant sets the limit to the certainty and order.

4 Time

If there is no perfect order, then there has to be a degree of disorder. However, disorder is just the opposite of order. In that sense, disorder is hardly a new concept. If something is not labeled as "order", then it has to be labeled as "disorder". Taken this way, order and disorder are complementary concepts. This complementarity follows from the idea that there has to be a perfect symmetry — the order and disorder should add up to a certain complete picture.

The concept of time breaks the paradoxical symmetry between order and disorder. Genuine uncertainty cannot be divorced from the concept of *time*. No static pattern can be considered as genuinely uncertain. No matter how random it may look, if it is static, it is there to be described in certain terms, to be copied, and even changed in a predictable way. What is uncertain is which of the possible options for future events will actually occur. Of course, effects of past uncertainty can be identified in the past events, but the past is determined. It is the present and the future where the grains of genuine uncertainty are. Once an event occurs, it becomes determined; it becomes static or dead — it belongs to the past.

4.1 Balance

There is a balance between order and uncertainty. They are antagonistic concepts that are both needed for a process to exist. A superior balance between these two antagonistic concepts results in a longer process. The balance is as fundamental for a long-lasting process as is the successful replication of a pattern for the dominance of that pattern in space.

The importance of the concept of balance can be illustrated by the evolution of complex plants and animals. A gene with high copying fidelity that is not balanced by uncertain mutations does not lead to a long-lasting evolutionary process. The total dominance of a gene with high copying fidelity would mean that all the available organic matter would be used for copies of that gene, ending the process of evolution. A gene with high copying fidelity, however, is very bad at adapting to the inevitable changes in the relevant conditions that ultimately halt the copying expansion of that gene. It is the genes with a better balance between copying fidelity and uncertain mutations that evolve into more sophisticated forms able to adapt and survive longer.

Expressed simply, a better-balanced process lasts longer. The replication is a simple mechanisms of gaining dominance for a pattern, but the balance between the replication and the uncertain mutations is a higher-level concept that can win even in the replication "game".

4.2 Random Events

Time without *events* is nothing. It may seem that a suitable mathematical definition of time is a continuous flow of something, but without events, it means nothing in reality. It is the events that link the concept of time to reality.

Complex things can be defined as events; for example, World War II can be called an event in history. Obviously, this kind of events has multileveled ingredients: battles, specific decisions of generals, deaths of individual soldiers, etc...

It is of interest to focus on the concept of an *elementary event*. If the granulation of time cannot continue in reality all the way to the mathematical abstraction of infinitesimal (dt), then the granulation of events cannot continue to an infinitely small event. Based on the uncertainty principle, the limit to a reduction of the length of an event is set by $\Delta t \Delta E \geq h$, where Δt is the uncertainty in the time taken by the event and ΔE is the uncertainty in the energy associated with the physical occurrence of the event. Clearly, the time taken for an elementary event to occur cannot be separated from the associated energy. It takes both time and energy for an event to occur. There are different possibilities in terms of how much energy and how long time are necessary. What is not possible is the product of the needed energy and time to be smaller than h.

Reducing an event to the uncertainties of its ingredients (ΔE and Δt) means that this event is a random event.

5 Creation with Natural Selection

The term *creation* is usually used to describe that a fairly complex thing is brought into existence: creation of a theory, creation of a piece of art, creation of a plant, creation of a human... These complex creations have multileveled structure that emerges after lengthy evolutionary processes. Again, the insightful question is about the smallest elements in this process, and specifically, the elements that the novel feature(s) originate from. Creation is something that requires time and energy, therefore, it cannot be granulated beyond the single *events*. For a novel creation, the event has to be something that is new, something that was not fully predictable the moment before the new creation occurred. Accordingly, we come to the idea that a *random event* is the smallest element, or the unit of creation.

Creation of a novel thing has to originate from a random event. If there is no uncertainty, then there is nothing new, nothing that was not determined before the event itself.

Random events can be put together to form a bigger piece of a newly created thing. For example, a monkey hitting randomly the keyboard of a word processor can create a lengthy piece of a genuinely new (unpredictable) text. Even better example is a monkey with a paint brush and a canvas. There can be no objection about the novelty of either the "poem" or the "picture" created by the monkey. The problem with these pieces of art is not with the novelty, the problem is that they are not likely to be successful. These novel creations will not do well through the inevitable *selection* mechanisms, so they will not get very far.

If "making a significant impact" is embodied in the definition of *creation*, then the creation becomes a *process* that involves *selection*. The greatness of a creation with an impact is due to the selection; the very origin of the new thing is as simple as a random event.

The simplest and the inevitable selection is known as the natural selection. Of course, the concept of *natural selection* was introduced by Charles Darwin as an essential component of the theory of evolution — the theory that explains the creation of plants, animals, and the humans. In the case of biological systems, the natural selection relates to the replication and the long-term survival of plant and animal species. A useful example is the creation of the commercial fruit varieties and domestic animals. The selection and repeated planting of a rare but tasty wild fruit leads to its dominance. Furthermore, new varieties are created by cross pollination (or cross breeding). The theory of evolution is usually presented as natural selection of random *mutations*, but a random shuffle of male and female genes in bisexual plants and animals is more recent and far more efficient method of creating superior novel plants and animals than the random-mutation method. With repeated cross pollination (or cross breeding) and selection, people have created abundance of plants and animals that suit their taste or need.

Being a great concept, the natural selection is also applicable at both lower and higher levels than the level of plants and animals. Crystals have been used as an example of low-level order. Randomly moving atoms may create many different patterns as they interact with one another, but if there is a well-balanced pattern[2] that lasts much longer, it is this pattern that will make a big impact. To lead to a large crystal rather than a stable molecule (such as H_2), this pattern should create favourable conditions for its replication. Once a number of atoms create a seed of this pattern, the seed acts as a kind of template for the creation of orderly crystal from the random motion of the

---

[2] *Minimum-energy cell* is the usual scientific terminology for the concept of *well-balanced pattern*.

supply atoms. Nonetheless, a repeated selection of a specific pattern is not to be confused with absolute determinism. Repeated selection of a pattern — replication — can create a highly-ordered and a relatively large crystal but it does not create a *perfect* crystal. The inevitable crystal defects (including surface termination) are so random that different pieces of the same type of crystal are not absolutely identical, although the differences can be ignored as irrelevant specifics for the purpose of suitable generalisations.

At the higher level of human ideas, a good example of natural selection is Thomas Edison's legendary failures to make the light bulb. Edison's famous spin is that he did not fail more than thousand times in his determination to make the light bulb but that he learned more than thousand ways how not to make it. In spite of Edison's knowledge spin, the huge number of attempts makes it possible to ignore the knowledge aspect — for a moment — and to use this as a superb example of a human invention largely based on random mutations of the central idea and the natural selection. A random change[3] creates a novel recipe that is tried, and if it does not work, the life of the invention (the novel recipe) is terminated. The random recipe that eventually worked was naturally selected as the creation that would be replicated and further improved as we continue to make the huge number of light bulbs.

6 Creation with Knowledge-Based Selection

Having distilled the concept of natural selection and its applicability at the level of human creations, Edison's knowledge element can be considered. To analyse the effect of the claim that he *learned how not to make bulbs*, imagine that the experiments were conducted by a different person who did not use the failures as knowledge. In the thousands of attempts, that person would have repeated already failed recipes, possibly many times. This "knowledge-less approach" is not absolutely hopeless. It is likely that a much longer perseverance would have been needed, but with that, the light bulb could have also been created. The very significant difference between the two approaches is in terms of efficiency. Importantly, this may not be just a quantitative difference, for example, if perseverance longer than a lifetime is needed. A mere record on failed recipes may be considered as a rudimentary level of *knowledge*, but when introduced into the selection process, it makes a dramatic impact.

The importance of knowledge for human beings and human society does not need a lengthy introduction. Knowledge is the core business of universities,

---

[3] A possible argument that Edison's mutations did not have the random element is contradicted by the mere fact that numerous attempts failed: if a recipe does not have the random element, then it is deterministic and it should not fail.

in two related ways: (1) knowledge dissemination (usually called teaching or *education*) and (2) knowledge creation (usually called *research*). Of course, education is also the core function of secondary schools, primary schools, and even kindergartens. Furthermore, a form of education or training is found in virtually every organization of people. Accordingly, it may seem that there are many different types of knowledge with different roles for different individuals and for different groups of people (societies). In fact it is more the other way around: people and societies become different because they develop different levels and type of knowledge.

The fundamental role of knowledge is very simple and common for everybody. It relates to the continuous need to make selections of available options, and obviously, it helps to make the more favourable selections. The phrase "selection of available options" is used because it is more descriptive, although it can be replaced by the more usual and the shorter term *decision*.

The following could be a helpful example. Imagine a distant predecessor of us humans having to make a selection whether to eat or not to eat a specific fruit. The necessity to make this kind of selections is unavoidable. Furthermore, making selections from the available options is the basic element of all our actions. Of course, these selections can be random. Assume that our distant relative made the selection to eat the fruit and that the fruit was fatally poisonous. This is an example of a random selection, and of course, very unlucky one. With appropriate knowledge, the selection would had been not to eat the fruit. It is the selections of possible options that create the future. In this example, the future was the death or life of our distant predecessor.

People, in particular managers, who religiously believe in *planning* may see a problem with the idea that the basic element of all our actions is simply *making selections from the available options*. Planning is, of course, at a much higher level than selections from randomly available options. To extend the example with the fruit, we need to assume a less distant predecessor who knows that seeds can be planted to produce fruits in the future. With this knowledge, the predecessor has the option of planting the seeds of a familiar and nice tasting fruit that was coincidentally found. If the option of growing own fruit is selected, a plan that is consistent with this selection would be devised. The plan could be quite straightforward: collect the seeds, plant them (select position), look after them (select methods, such as watering, protection from birds, etc.), and collect the newly grown fruit. Selections, such as watering as a method of looking after growing plants, can be planned. Nonetheless, the selections have to be actually made not just planned. Imagine a lazy plant grower who selects not to water the plants a day by day and a week by week until the plants die. It was the selections whether to water or not to water that mattered, not the mere plan to water them.

Clearly, planning does not remove the need to make selections from the available options. What the planning does is to guide the creation of new options. A success with the fruit-growing plan would open many options that go with the extra time that does not need to be spent on random search for the food[4]. A possible abundance of fruit would also create many new options.

It may seem that the best possible planning is to develop the plans to the level of precise algorithms that can run as smoothly and as reliably as a mathematical algorithm in a perfect computer. This kind of plan is unachievable, because it is impossible to predict the inevitable random events and plan for them. For example, a total reliance on a seemingly perfect plan for growing own fruit is not the smartest plan. An unpredictable storm may wipe out all the grown fruit, and without a viable alternative option for the needed food, not only would the plan be dead but everybody who depended on this plan may also die. The smartest plan is not the perfect plan but a plan that involves the most adequate balance between an orderly algorithm and uncertainty.

Knowledge is needed for a good plan. If the knowledge is not there, then it is better to create it than to continuously rely on random selections. A simple plan to create knowledge for the first time could be based on natural selection, or in other words, on a real-life *experiment*. How did our predecessors get to know which fruits are poisonous? A possibility is that there were "smart managers" who convinced the bravest members of the group to try the unknown fruit, with the following plan: if everything goes well, then both the manager's knowledge and the manager's status rises significantly, opening new options; if the attempt does not go well, the manager could still use the newly acquired knowledge even if it becomes necessary to move to another group.

6.1 Memorized Knowledge

All the previous examples of knowledge can be classified as *memorized knowledge*. Memorized facts enable the most straightforward and the fastest decisions (selections of available options). This is the lowest level of knowledge that evolved very early (even primitive animals are able to make favourable decisions based on memorized facts). In addition to genetic inheritance, animals and subsequently humans have developed ways of communicating the acquired knowledge. Speech is relatively recent and quite sophisticated form of knowledge communication; older forms of knowledge dissemination are based on drawings, use of simple symbols, and body language. Therefore, education

---

[4] A smart selection for at least a part of this extra time would be to upgrade the knowledge so that even better plans could be devised.

is older and more fundamental than the speech and even older than humanity. This means that humanity is impossible without education, which further means that a decline in the quality of education inevitably leads to a decline in the society in general and to a decline in the quality of life conditions for its individuals. The big question here is about the essential characteristics of *quality* education. At this stage, it is safe to conclude that an education system that is focused on memorizing facts has to be regarded as the lowest-level education. This does not mean low quality if this system is used to teach children the symbols of the alphabet, but it certainly means low quality if it is applied at the university level.

The memory in the brain of animals and humans is quite limited. A demonstration of this memory limit is the fact that people need to forget many facts from the past to be able to memorize facts that are relevant for current decisions. The discovery that the brain's memory can be extended by external records enabled the development of memorized knowledge to much higher levels than the simple memorizing of facts. Put simply, it was no longer necessary to remember the facts themselves; it became sufficient to remember where to find and how to use the records about the known facts. This approach becomes especially efficient with the development of hierarchical structures (order) of the established facts. The simplest form of structured knowledge can be illustrated by the example where we need to find a fact about certain fruit, say oranges. We could go to the section of a library that has the plant-related books, find a book on fruits, and look through the citrus chapter. Another form of structured knowledge was enabled by the development of symbolic mathematical equations to replace tables of experimentally obtained numbers. An example could be the experimental observation that the current through a resistor depends linearly on the voltage across the resistor. Instead of recording tables of current versus voltage values for different resistors, Ohm was able to record this experimental fact in a very compact form: $I = V/R$, where V can be replaced by a voltage value and R by the resistance to obtain the corresponding current, labeled by I. Yet another form of structured knowledge relates to recipes or algorithms as efficient records of established facts about different processes. An example here is, of course, Edison's recipe for making light bulbs. An analogous and a widespread example is the cookbooks that provide step by step instructions how to make cakes and other dishes.

The memorized knowledge, both externally and in human brains, is very useful for making the necessary selections in the process of creating copies of previously invented things. To make a well-known cake ordered by some customers, a cook uses the adequate recipe (memorized knowledge) to decide that x grams of flour should be measured, followed by y grams of sugar, z

grams of milk, etc., then to decide to mix the ingredients for certain time, to bake it at certain temperature, and so on. If the cook suddenly realizes that salt instead of sugar was mixed by omission, which is a random fluctuation, the decision would be to throw the mix away and start again. It would have been an original creation, but the memorized knowledge is used to conclude that this is not a useful creation because it is not the cake that the customers want.

Creating copies of previously invented things is the concept of order created by replication. This concept has been very successfully utilized by the mass-production industry to provide us with very affordable yet very sophisticated creations, ranging from cars and computers, through movies and books, to training courses. A lot of people are employed by this industry, so they have to be educated. This kind of education is usually achieved through apprenticeships or by so called vocational training.

Replicating stuff hour by hour, day in and day out, is not the best balance between order and uncertainty for humans. This is why hobbies and activities that involve uncertainty, such as games (chess, sport games, fishing, etc.), reading novels and watching movies with unexpected twists, and jokes are so popular. This is why humans are so keen on developing machines to perform the replicating work for them.

6.2 Thinking

Although very useful, the effects of memorized knowledge are limited to replications. This is because the very role of the memorized knowledge is to select creations that are known to work. The example with the original cake, created by the random replacement of sugar with salt, shows this. Because it is new — it does not belong to the existing knowledge on verified recipes — it is uncertain whether it would work with the customers, so it is thrown away. The alternative to this decision is not to disregard the new recipe but to give it a chance. This can happen through natural selection, in analogy with Edison's recipes for the light bulb. For a natural selection in this case, the novel cake recipe can be tried with the customers. The chance that they will like it more than many existing cakes cannot be excluded, and if they do, the new creation would be successful and a rewarding replication of the new recipe would begin.

Although the method of creation by natural selection works, it is very resource consuming (both in terms of time and materials/energy). There are far too many possible ingredients and many more possible random combinations, and most of them will not work. Certainly, the cheapest and the fastest decision that can be made about a new creation is to dump it. The cheapest option, however, may not be the most rewarding option. It may be far more rewarding to extend the selection process beyond mere comparisons to memorized knowledge. In the example with the novel cake, *thinking* may indeed lead to the conclusion that it is very unlikely the customers would accept a salty cake for desert. With this conclusion, no natural selection would be attempted, which saves money, time, and embarrassment. However, the thinking process itself is affected by random events. So, imagine that a random event causes the thinker to start comparing the novel *cake* with established *pie* recipes, and to discover in this way that the novel cake recipe is an interesting variation of a very popular pie. With this, the novel recipe begins to look like a real invention. Further thinking through random variations of several spices may produce further development of the idea, and eventually, a solid plan for experiment (natural selection).

A mental process for selection of ideas that goes beyond a mere comparison to memorized knowledge is *thinking*. By not suppressing every single novel idea, just because it is not a part of the memorized knowledge, thinking allows inventions to occur. On the other hand, by classifying most of the novel ideas as unworkable (as most random-based ideas would be), the thinking-based selection is far more efficient than the natural selection.

As a higher-level concept than memorized knowledge, thinking is impossible without memorized knowledge. Naturally, the transition between selections based on memorized knowledge and selections made after a thinking process is very subtle. Most thinking processes are simple and short, needed to estimate an unknown situation and to make a decision for action. These thinking processes are still original, given that they deal with initially unknown situation, but they are not aimed at long-term plans and great inventions; they are aimed at short-term actions and survival.

Random ideas and the thinking-based selection process lead to analogy with the evolution (Darwinian process). William Calvin posted the following sentence on his web page (http://williamcalvin.com/): "I tend to think that the fancier mental processes (language, planning, music, logic) utilize a form of Darwinian process that operates in milliseconds to minutes." Given that the evolution (Darwinian process) is a great concept, the postulate that great concepts do not remain hidden at lower levels leads directly to the conclusion that evolutionary processes have to appear at the higher level of human ideas and thinking. Thus, the logic presented here is in a full agreement with Calvin's sentence. I am not so sure about the specific model with hexagons that is proposed in his publications[4]-[6], but this could be due to my poor knowledge of brain anatomy.

There is probably a widespread belief that no random events could occur in the thinking brain, which would be inconsistent with the concept of evolutionary processes at the level of ideas and thinking. In fact, there is plenty of room for random events as the brain signals travel through the maze of neurons, in particular at the synapses (the junctions between neurons). Dreams can provide an indication of the appearance of complex random ideas in the lower (subconscious) levels of the brain activities. This example also indicates that the brain has a subconscious mechanism for prioritizing the ideas that are to be thoroughly considered at the thinking level.

Just as in the case of selections based on memorized knowledge, the thinking process is greatly enhanced by structured knowledge. This is in a full agreement with yet another sentence from Calvin's web page: "To keep a half-dozen concepts from blending together like a summer drink, you need some mental structuring." It is the type and the degree of structure (order) in the knowledge that is memorized in someone's brain that makes the decisive impact on the thinking process.

Take the following example of two pieces of knowledge: (1) the already mentioned Ohm's law ($I = V/R$) and (2) the energy-conservation law. Imagine two researchers discussing problems with their experimental data. The problem of one of them is that I–V measurements made with a semiconductor resistor do not agree with the Ohm's law. The problem of the other researcher is analogous: analysing calorimetric measurements, this researcher finds out that the experimental data do not agree with the energy-conservation principle. The decision what to do in each case depends on the particular structure of the knowledge of the thinker making the decision. A thinker without a knowledge structure regarding these two laws, for example a thinker who considers both laws as fundamental truths, will conclude that both sets of experiments are wrong. Different conclusions are drawn by a thinker with a structured knowledge. The structured knowledge in this case could be as follows. The energy-conservation law is a fundamental principle, whereas the linear dependence of current on voltage is just a widespread first-order (linear model) approximation. This leads to a different thinking process: there is an error with either the measurement or the analysis of the energy related experiment, whereas the result from the I–V measurements deserves a real attention.

Decision making in complex situations critically depends on the structured knowledge and its application. Complex situations, such as those typically faced by managers, may involve conflicting factors and most likely involve uncertainty about future factors, including related future decisions of others. A possible approach is to invest time and resources into a natural-selection process to minimize the uncertainty about the future. For example, if the uncertainty relates to technical issues, the investment into proper experiments is likely to remove it. Another example is when the uncertainty relates to market issues, in which case a proper market survey may do the job. Although useful, a systematic natural selection of ideas can be prohibitively expensive in terms of both time and resources. Another option is to submit the situation to natural selection by not making a timely decision. Yet another option is to make a random decision: this type of decision may seem inevitable or the decision maker may not be realizing that the decision is essentially random. Despite of all these options in a complex situation, it may be possible to reconsider the situation at a higher level of the structured knowledge — to take the big-picture view. This enables to ignore many factors that become irrelevant specifics at the higher level. Also, the short-term uncertainties give the ground to more enduring "principles". The higher the level, the simpler the situation becomes, enabling an easy decision. Of course, this type of decision may work against some short-term or specific wishes, but in long term, it is the best decision.

Any attempt to create a novel thing is a highly complex situation, at least because it inevitably involves a high-degree of uncertainty about future factors. In these situations, an adequate and a highly structured knowledge is necessary to be able to make the proper big-picture decisions that create the future. In the absence of this knowledge, the decisions are random and the future is created by natural selection.

Structured knowledge is a high-level order that can be considered as a kind of branching cause-and-effect links. A common cause can be found for many seemingly different effects when irrelevant specifics are ignored to move to a more general level. *Finding a cause* is equivalent to answering a set of questions that are related to the observed effect(s). In many instances, the bottom-line question is a *why* question. If an answer to the questions cannot be found in the memorized knowledge, then they become a valuable learning tool. It will be necessary to go through a thinking process to find satisfactory answers, and this thinking process may need to be supported by many other activities. When the answers are found, knowledge at a more general node in the cause-and-effect tree is created. This is referred to as *understanding*. Importantly, as the cause of a single effect is understood, it becomes obvious that many other effects are related to the same cause. This opens new options for application of the newly gained structured knowledge. The most effective plans for future creations are those that are built on sound understanding of the relevant cause-and-effect links.

A newly created structured knowledge is a result of individual thinking process and can be original to the level that it enables decisions that are not obvious to others. Creations emerging from understanding that is not obvious, or from plans based on such understanding, are original creations — these are the knowledge-based inventions.

Questions that could not be answered by the memorized knowledge were referred to as a learning tool. However, if the questions trigger an invention, then they are more than a learning tool; these questions end up being a knowledge-creation (research) tool. No question can be classified with certainty as a research tool. When a thinking process is initiated by a question, it is impossible to predict with certainty that it will lead to a new knowledge rather than to a fact or principle that is already known by others. This inseparable link between learning and inventing is referred to as teaching and research nexus in the University jargon. The nexus is in the thinking process that is essential for both the learning of structured knowledge and the knowledge-based inventions.

Cause-end-effect chains are not infinite. As in the Schrödinger-cat example, every cause-and-event chain begins with a random event. Consequently, the use of structured knowledge for great plans and great inventions is limited. It cannot be extended to enable perfect plans that rely on perfect predictions of the future. A plan to create the future that is based on adequate balance between structured knowledge and uncertainty is the best possible plan.

Epilogue

The thinking presented in these lecture notes is based on the central postulate that the effects of great concepts are not limited to microscopic levels but flourish with the more sophisticated opportunities at higher levels of complexity. The main aim has been to demonstrate that a consistent and logical model, accounting for as elusive concepts as creation and knowledge, can be based on this postulate. The most controversial element in the model would be the concept of uncertainty as the fundamental limit of any cause-and-event chain and its links through the quite obvious uncertainty and randomness at higher levels of complexity to the fundamental role of random events at the level of original creations by humans.

Although the material presented in these notes has not been subject to a peer review, and in that sense this is not a formal publication, the copyright is defined and it belongs to the author. If you find the material provocative, or just wish to express your opposing view to someone, it will be very much appreciated if you contact me with your comments. Questions and positive comments will also be very much appreciated. There is no intention to ignore questions, and some are obvious: how does the presented thinking model relate to consciousness?

References

[1] R. Dawkins, *The Selfish Gene*, 2nd. ed., Oxford University Press, Oxford, 1990.

[2] J. Gribbin, *In Search of Schrödinger's Cat: Quantum Physics and Reality*, Black Swan, London, 1991.

[3] R. Feynman, R. Leighton, and M. Sands, *The Feynman Lectures on Physics*, vol. III, Addison Wesley, Reading, MA, 1965.

[4] W. Calvin, "Competing for Consciousness: A Darwinian Mechanism at an Appropriate Level of Explanation", *Journal of Consciousness Studies*, vol. 5, pp. 389-404 (1998).

[5] W. Calvin, *How Brains Think*, Basic Books, New York, 1996.

[6] W. Calvin, *The Cerebral Code: Thinking a Thought in the Mosaics of the Mind*, MIT Press, Cambridge, MA, 1996.