

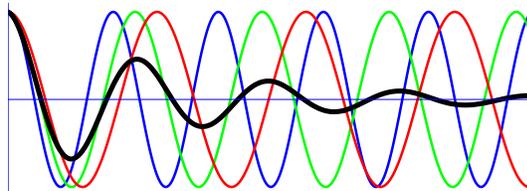
An introduction to the basics of dephasing

Florian Marquardt
(Florian.Marquardt@unibas.ch, University of Basel, Switzerland)

10th August 2001

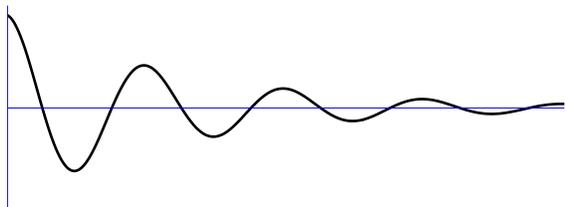
Abstract

I offer a simple introduction to the basics of dephasing (or “decoherence”), i.e. the process by which quantum-mechanical interference is destroyed. Emphasis is placed on the physical principles and concrete examples, rather than the formalism.



Oscillations fading away

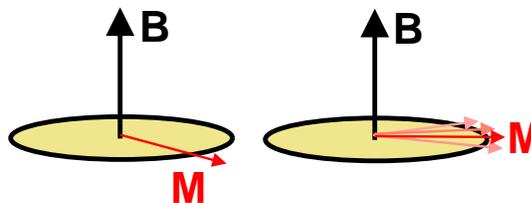
Oscillations are ubiquitous in physics, and so is their decay, shown in the picture below:



Maybe this curve just represents the amplitude of a pendulum, which loses its energy due to friction and finally comes to a rest. Then it would have little to do with the phenomenon which we call “dephasing” and which this short introduction is about. But maybe the graph really shows a *superposition* of many oscillatory curves, each one of them having a slightly different frequency - see the next picture (where the frequencies are distinguished by appropriate colors). Initially all the oscillations start with the same phase, but, with time progressing, they “get out of phase”, so their average value decays.

Spins in a spatially varying field

An example for this kind of behaviour is an ensemble of spins in an inhomogeneous magnetic field (pictured below), when the field strength differs between the places of the spins. Since the field controls the precession frequency, the average magnetization shows decaying oscillations, just like it has been explained above. This effect is important in nuclear and electronic spin resonance experiments and is one of the first examples where “dephasing” was analyzed for a quantum-mechanical two-level system, the spin. To observe it, one first applies a Rabi-pulse of an (additional) oscillating magnetic field in such a way that all the spins are flipped from their initial direction (pointing along the external field) into the plane perpendicular to it. Then the components of the magnetization within that plane will show the decaying pattern displayed above.



The spin is a particularly nice example to illustrate the effects of dephasing. This is because a spin lying in the xy plane is described by a quantum-mechanical superposition of the states “spin up” and “spin down” along the z -direction, both having equal weights (probabilities). The phase-factor $e^{i\varphi}$ which appears in this superposition takes on a direct geometrical meaning, since the phase φ is equal to the angle φ of the spin in the xy plane. The time-rate of change in φ , i.e. the precession frequency $\dot{\varphi}$, is given by the energy difference between the up and down states, which is proportional to the applied magnetic field. Of course, this kind of picture can be applied to any two-level system, even if the two states have a rather abstract meaning.

Fluctuations in time

Besides fluctuations in *space*, another source of dephasing in an ensemble of spins may be fluctuations of the magnetic field in *time*. In such a case, the phase of the precessing spin advances with a rate which is constant only on the average. Sometimes the phase will grow faster than usual, sometimes it will lag behind, giving rise to a fluctuating time-evolution of the magnetization itself. Again, a decaying oscillatory pattern will result if one averages this over many spins, each of which is subject to different time-dependent fluctuations. Alternatively, one could run an experiment many times on a single system and average over the outcomes of many runs.

At this point, it is interesting to note that not all kinds of fluctuations will lead to a complete decay of the oscillations. Maybe positive fluctuations of the magnetic field (increasing its value above average) are followed by fluctuations pointing in the other direction, which tend to cancel the effect on the precessing spin. Then, although the phase shows some jitter around its average value, the deviations from that value cannot increase with time. A complete decay can result only if positive and negative fluctuations do not cancel. In that case, the phase may perform a kind of random walk, deviating ever more strongly from its average value. However, it is not uncommon for the fluctuations to be of the first, “harmless” variety: For example, the thermal fluctuations of the magnetic field in free space will be exactly of this type!

Fluctuations which merely change the phase

Now, you should not start to believe that a spin will not be dephased by the thermal fluctuations of the magnetic field. This impression may arise here merely because, up to now, we have assumed that only the “ z -component” of the magnetic field fluctuates, which is the one that controls the precession frequency, via the energy difference between “up” and “down”. Such a sort of external influence, which only changes the *energies* of different quantum-mechanical states, is called “diagonal coupling”, because the quantum-mechanical operator representing the action of the external field is diagonal in the eigenbasis of the system to be dephased, i.e. the spin. It is the simplest kind of coupling (at least for calculations).

Fluctuations that can flip the spin

However, the fluctuating magnetic field usually will have components that point into directions other than the “special” z -axis which is selected by the strong, externally applied static field. These other components can even tip the spin out of the xy -plane. And, furthermore, for them the argument about the cancellation of “positive and negative” fluctuations does not work anymore, so that, indeed, the thermally fluctuating magnetic field will randomize the phase of the spin *and* its angle with respect to the z -axis. This coupling is called “nondiagonal” (as you may have guessed). It can flip an “up” spin into a “down” spin and vice versa, thereby changing the occupation probabilities of “up” and “down”. Since these states have different energies (when they are in the static external field), an energy exchange is involved that takes place between the spin (the “system”) and the fluctuating field, which is often called a “bath” in the general context of dephasing. The name “bath” derives, of course, from the heat bath considered in thermodynamics, which is held at a given temperature and drives a small system into equilibrium when coupled to it. We will come back later to discuss the distinction between dephasing and the more common effects of heating and energy dissipation that are also produced by a heat bath.

Fluctuations are harmless when too fast or too slow

Of course, the strength of the fluctuations is important in determining how fast the spin is dephased. But even more important than the strength is the frequency spectrum. If the fluctuations are very fast, they tend to average out and their effect is relatively small: The dephasing time will be long. In nuclear magnetic resonance, this applies especially to molecules in a liquid. Since their surrounding changes when they move around, the magnetic field acting on a nuclear spin inside the molecule fluctuates rather rapidly. That, in turn, increases the dephasing time and therefore leads to a sharpening of the transition line in the NMR spectrum, compared to the situation in a solid. This effect is called “**motional narrowing**”. One can understand it better if one realizes that the phase performs a kind of random walk due to the fluctuating field. The total average distance travelled by a random walker gets smaller when the steps themselves get smaller both in duration and in size. This corresponds to the rapid fluctuations of the magnetic field acting on a spin in a liquid.

In the case of NMR, the magnetic field produced by the spin of a neighboring proton corresponds to a frequency shift of about a few tens of kilohertz. Therefore, if there were no fluctuations in the positions and orientation of neighboring spins, the dephasing time would be less than a millisecond. Usually, however, it is rather on the order of tens of milliseconds up to a second, owing to this averaging phenomenon.

On the other hand, dephasing may also be suppressed if the fluctuations are very *slow*. This is because the effect of the fluctuating field (in the x or y directions) depends on the phase of the precessing spin. If the spin’s precession frequency is much larger than the typical frequencies of the random field, the influence of the latter is diminished - again because of the averaging argument. If you tried to apply this last argument to the fluctuations in the z-direction (the “diagonal coupling”), you would fail: There, the influence of the field does not depend on the current phase of the spin. Because of this, it will not cancel out even if the spin precesses very fast. In fact, for the diagonal coupling, the low frequency fluctuations are the most harmful: they can never be “too slow”.

Thus, this rather simple situation makes it clear why the

frequency spectrum of the bath matters so much in dephasing.

Spontaneous emission of energy into the “bath”: the example of the atom

Up to now, the fluctuating field has been treated as if it were imposed on the spin as a kind of classical (but random) external force. What has been neglected is the back-reaction of the system (spin) onto the bath. This *is* important, since it is the basis of energy dissipation - and it also leads to dephasing, as we will see. In discussing this, let us take an atom as an example, instead of the spin. Although even a fluctuating classical electromagnetic field is able to de-excite the atom (“induced emission”), it is equally likely to pump energy into it. Therefore, in the long run the atom would have an equal population of ground state and excited states, which corresponds to equilibrium at *infinite* temperature. Only because there is the extra energy dissipation provided by spontaneous emission of photons, will there be a difference in the populations of the atom’s energy levels. This is especially important at low temperatures - and for an atom, room temperature is practically zero, when we regard the transitions between the ground state and excited levels which have a few electron volts of energy! Then, energy can only be *emitted* by the atom into the field. This falls into the category of a “nondiagonal” coupling between bath and system - i.e. one which changes the populations of the various energy levels by “flipping” the state.

When does it show up as a kind of dephasing? Imagine shining a strong laser beam onto the atom. This produces coherent “Rabi oscillations”, where the atom continuously cycles between ground and excited state. Their frequency will get larger if the beam is more intense. If you do it with an ensemble of atoms, you can record these oscillations by detecting the photons which are spontaneously emitted by the atoms in the upper level (“resonance fluorescence”). Immediately after switching on the beam, you will see clean oscillations in the brightness of the resonance fluorescence. After a while (typically some nanoseconds), the oscillations fade away and only the average brightness level is observed. This is because every spontaneous emission resets the atom to the ground state.

Since this happens at random points in time, the oscillations in the populations of the two atomic levels get out of phase for the different atoms in the ensemble. By the way, in this example spontaneous emission is both the source of dephasing and the means of observing the coherent oscillations.

The role of vacuum fluctuations

How does the “frequency spectrum of fluctuations” come into play here? After all, there are no (thermal) fluctuations at zero temperature. However, both the effects of thermal fluctuations and those of spontaneous emission depend on the density of field modes of the electromagnetic field. In the case of spontaneous emission, evaluating the mode density at the transition frequency directly determines the decay rate, since it counts the number of modes a photon can be emitted into. Typical spontaneous decay lifetimes are on the order of nanoseconds. This is about a million times longer than the oscillation period 10^{-15} s of an atom which is in a superposition of ground state and excited levels. Therefore, one could have a million cycles of coherent oscillations in an observable like the atom’s dipole moment, after exciting the atom into such a superposition. (Of course, for a single atom, you cannot measure this in a single run! Every measurement fixes the state and destroys the coherent evolution. That’s why we talked about an *ensemble* of atoms in the example of the Rabi oscillations discussed above.) In the case of thermal fluctuations (induced absorption and emission), the mode density still has to be multiplied by the average number n of thermal photons in such a field mode, in order to find the decay rate of an atomic level.

In the general formula for the transition rates, the effect of spontaneous emission is then just to replace “ n ” by “ $n+1$ ” in the emission rate. This has suggested an interpretation of spontaneous emission as being a kind of induced emission due to the “vacuum fluctuations” of the electromagnetic field - i.e. due to the fact that the fluctuations of electric and magnetic field values do not vanish even in the ground state of the field. That, again, is merely a consequence of the finite extent of the ground state wavefunction of a harmonic oscillator - the electromagnetic field being merely a collection of independent normal mode oscillators. However, this view turns out to be a bit oversim-

plified, since it leaves unanswered the question why these vacuum fluctuations do not also *excite* the atom from its ground state. A consistent way of looking at these things is to say that spontaneous emission from an excited level is partly due to vacuum fluctuations and partly due to “radiation reaction” (the damping force acting on any radiating charge). For an excited state, these act into the same direction, while they cancel for the ground state (leading to a balance of fluctuations and dissipation).

Therefore, although the vacuum fluctuations are important in decay and dephasing, we have to take care not to treat them on the same footing as the “usual” thermal (or arbitrary classical) fluctuations. They cannot be separated from the corresponding dissipative processes (like radiation reaction, in this example), because the latter necessarily are of the same strength. Their main effect will be to lead to transitions downwards in energy (spontaneous emission), but they will never excite the atom. Of course, this is just a basic consequence of energy conservation: While a photon can be emitted into the field, there are no photons to be absorbed out of the field, at zero temperature. This is not to say that there are not also other, more subtle effects of the vacuum fluctuations, like for example energy shifts, smearing of wave functions and a change of the effective electron mass.

The frequency spectrum of electromagnetic vacuum fluctuations

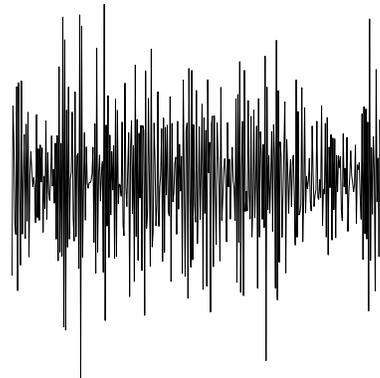
Let us come back to the question of the fluctuation spectrum in the case of the atom. The behaviour of the spectrum of the electric field fluctuations acting onto the atom can be obtained in two steps: For the vacuum in three dimensional space, the relevant density of electromagnetic modes grows like ω^2 with increasing frequency ω . Since each field mode carries an energy of $\hbar\omega/2$ (in its ground state), and since the energy density goes like the square of the electric field, the frequency spectrum of the electric field vacuum fluctuations grows very fast, like ω^3 . This means that the spontaneous emission rate also grows like ω^3 - if all other things are kept equal (in particular the dipole matrix element of the transition). This fast growth with frequency has the consequence that transitions with a *small* energy difference (frequency) will be very long lived. If you want to preserve coherence for a long time,

this is good news. Of course, you have to pay a price for it: All the dynamics taking place in a system with such a low transition frequency usually will be very slow as well. But since ω^3 falls off faster than ω at low frequencies, you might still be better off with a low-frequency system, if you want to observe coherent oscillations for many periods.

In the case of atoms, this is relevant for Rydberg atoms, i.e. atoms which have been excited into a high-lying level with a large quantum number (for example $n=30$). Since the electron is so far away from the nucleus, its velocity and the frequency of its orbit are very low, which means small transition frequencies to neighboring levels (they drop like $1/n^3$). Somewhat unfortunately, the dipole moment is necessarily very large, due to the large radius of this “outer” orbit. However, combining all the numbers still leads to a relatively long typical lifetime of Rydberg states, which is around a millisecond for $n=30$.

Although other baths will have different spectra, the decay of the spectrum towards low frequencies is a common feature. Therefore, as a rule of thumb, the general trend is indeed towards longer dephasing times (better coherence properties) when one goes to systems with slower dynamics (lower frequencies).

What is the actual **magnitude of the electromagnetic vacuum fluctuations**? If one wants to answer this question, one first has to define the frequency range over which the fluctuations should be taken into account. Taking the electron’s Compton frequency of about $10^{21} Hz$ (equivalent to its rest mass) as a (slightly arbitrary) upper cutoff, the strength of the field fluctuations is about *ten million times larger* than the electric field that binds the electron to the nucleus!! However, this does not lead to a catastrophe, because most of it is due to fluctuations at very high frequencies, which are much larger than the typical orbital frequency of the electron in the atom. They just lead to a “jitter” of the electron’s position, by about 1/1000th of the electron’s orbital radius. This smears the electron’s wavefunction a little bit, which increases the average distance to the nucleus for the s-orbitals and therefore shifts their energies to somewhat higher values (by around a Gigahertz, i.e. one millionth of the typical transition frequencies). The effect is the famous “Lamb shift”.



The vacuum fluctuations of the electric field.

If one wants to estimate the *relaxation* of an electron in, say, the first excited state of the atom, one should rather take into account only those fluctuations which are in resonance with the transition frequency on the order of $10^{15} Hz$. Of course, they do not have to be *exactly* in resonance. It suffices if they can extract energy from / feed energy into the atom resonantly for about the lifetime of the excited state, which is on the order of a nanosecond. (Actually, if you want to calculate the lifetime by estimating the strength of the fluctuations, this is a self-consistency problem). So we need to know the magnitude of electric field fluctuations at the transition frequency, with a frequency spread of a Gigahertz. It turns out to be roughly $10^4 V/m$, which is about ten million times *smaller* than the electric field binding the electron to the nucleus. And therefore it does not come as a surprise that, in order for this relatively small fluctuating field to have an appreciable effect on the electron in the atom, it has to act resonantly for a few million cycles - which is just the lifetime of the excited state, about a nanosecond. (Nature is self-consistent, after all.)



The fluctuations remaining after filtering out only a narrow frequency interval around the atom’s transition frequency.

By the way: Had we allowed for a frequency spread which is four times larger, the magnitude of the fluctuations would only have grown by a factor of two. This is just the old story that the spread in a sum of independent random variables (here: the frequency components of the electric field) grows like the square root of the number of variables added. And that behaviour is important here,

otherwise we would not get a fixed and finite lifetime (because having a smaller lifetime would increase the fluctuations would decrease the lifetime...)!

Changing the fluctuation spectrum

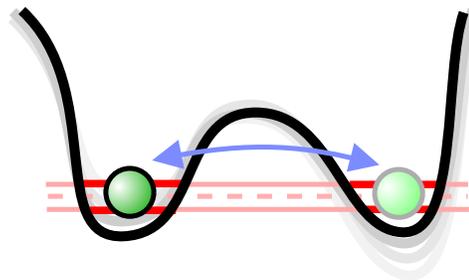
Finally, one might ask whether the electromagnetic field fluctuations can also be “too fast” or “too slow” for dephasing the atom. Now, their spectrum in free space is certainly fixed, but as soon as the boundary conditions are different, the modes of the electromagnetic field change. This is particularly important if the atom is enclosed in a small box, of dimensions comparable to the relevant wavelength - which is actually done in the experimental field of “cavity Quantum Electrodynamics”. Then, if there is a mismatch between the transition frequency and the available mode frequencies, spontaneous decay is indeed suppressed. However, even if the frequency of a mode is tuned to resonance with the atom, spontaneous decay is no longer an irreversible process as it had been in free space. The energy is exchanged continuously between atom and mode, in a kind of “vacuum Rabi oscillation”. Physically, one can simply imagine that the atom reabsorbs a photon which it had emitted previously but which has been reflected off the walls of the cavity. Thus, another important feature of a “bath” becomes apparent: it should lead to irreversible processes, which is (strictly speaking) only possible if it has a continuum of excitation frequencies - unlike the discrete spectrum of the field in a cavity.

Other examples of two-level systems

Once you know the story of dephasing for a two-level system like the spin or an atom, in principle you have understood it for every other two-level system as well. In practice, it still pays off to take a look at all those other examples. Although the two-level dynamics is of course always the same (“flipping” the “spin”, inducing Rabi oscillations and so on), the physical features of the various baths that are responsible for dephasing vary widely.

An important example are tunneling systems, where a particle can sit at the two minima of a double-well potential

and tunnel between them. Now “up” and “down” get replaced by “left” and “right” (see picture).



A double-well potential with a tunnel barrier: ground and first excited state (full lines) are separated by the tunnel splitting. The potential energies of left and right well may fluctuate (as indicated), leading to dephasing.

The particle can be an ion, a proton, an atom, a muon or an electron. The two wells of the double-well can be two energetically preferred locations in a crystal or a glass, which are close to each other. Maybe they even have the same energy (likely in the crystal, due to symmetry). The wells could also be the two possible positions of the nitrogen atom in the ammoniac NH_3 molecule - a classic textbook example of a two-level system. On the other hand, the double-well system also occurs in more abstract settings: for example for the two possible directions of current flow in a small superconducting SQUID ring or in the description of a chemical reaction.

Other examples of baths

In the solid, the “bath” may be the lattice vibrations that produce distortions of the charge density and therefore time-varying electric fields. In a metal, the conduction electrons are another bath. The fluctuation spectrum of the electric field then is determined (at low frequencies) by the Nyquist noise which is due to the random motion of electrons in the metal. On the other hand, in a superconducting metal, the electronic source of noise is strongly suppressed at low frequencies and vanishes exponentially at zero temperature. A case like the Fermi liquid 3He is in-between: although it is clean (and therefore does not

show Nyquist noise), the fluctuations acting on a particle immersed in the liquid are still stronger than they would be in a superconductor (or, for this case, in a superfluid). In general, electronic fluctuations are more important at low temperatures and low frequencies, while the lattice vibrations dominate at higher temperatures. An isolator, of course, does not show electronic fluctuations (at low frequencies), since there are no conduction electrons moving around. Therefore, coherence is maintained even better than in a superconductor. Compared with these sources of noise, fluctuations of the vacuum electromagnetic field are generally tiny (as we will show later in the example of Nyquist noise).

Experimental results on tunneling systems in solids

In some slightly disordered crystals, there are pairs of equivalent sites between which a particle can tunnel. For example, experiments have been performed using protons or positive muons. A typical value for the tunneling frequencies encountered there is 10^{11} Hz (for protons in niobium). At low temperatures, the interaction of the tunneling particle with the surrounding electrons is the most important dephasing mechanism. The decoherence rate grows linearly with temperature and, for about ten Kelvin, becomes roughly the same as the tunneling rate. Therefore, above ten Kelvin the oscillatory motion is overdamped - after starting the particle in one well, it will not even perform one cycle of coherent oscillations between the two wells. On the other hand, below nine Kelvin, niobium becomes superconducting. Then, the low-frequency fluctuations of the electrons are strongly suppressed and so is the dephasing rate.

If there is a whole lattice of equivalent sites for the tunneling particle, it may delocalize and move throughout the whole crystal. This has been observed for positive muons moving in a crystal of aluminium. However, usually this motion is not coherent but diffusive, since the phase coherence is lost after each tunneling event. At high temperatures, the tunneling is thermally activated (lattice vibrations helping the particle over the barrier) and therefore the rate increases with increasing temperature. On the other hand, at low temperatures the tunneling is quantum-mechanical and therefore its rate increases with *decreas-*

ing temperature, since then dephasing is weakened. In that case, one speaks of “quantum diffusion”.

In order to go to the regime of coherent delocalized motion, one has to suppress the dephasing rate much below the tunneling rate. This has been achieved for muons using the superconducting state of aluminium (at less than one Kelvin). There, dephasing rates of less than one Megahertz were inferred, compared with a jump rate of 10^8 Hz. In an insulator, typical dephasing rates for comparable quantum diffusion experiments are as low or even lower.

“Quantum Brownian motion”

Up to now, we have been occupied with two-level systems. What about dephasing in other simple quantum systems, like, for example, a free particle or a harmonic oscillator?

Of course, the multitude of possibilities for the bath is greatly increased, because now not only the frequency spectrum but also the details of the spatial coupling can vary: Is the fluctuating force homogeneous in space or not? If not, what are the typical wavelengths of the fluctuating field?

In order to get a first insight into the problem, one considers the simplest situation, which is that of a homogeneous force. Physically, this can be justified at least if the potential restricts the particle’s motion to a finite region of space. Then, if the relevant wavelengths of the field are much longer than the extent of this region, the force may be treated as essentially constant across space. This is routinely assumed for the atom, where this procedure is called the “dipole approximation”, because, after the approximation, only the atom’s electric dipole moment couples to the electric field.

Still, there remains the question of the fluctuation spectrum. One may have a specific model for the bath in a certain physical situation. On the other hand, one may be interested in generic features only. Since coupling to a bath leads to dissipation (friction) and physically the most important type of friction force is just proportional to the velocity, one may try a bath spectrum that will reproduce this behaviour in the classical limit. Thus, what

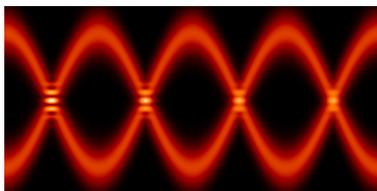
one asks for, is the quantum version of classical Brownian motion! The fluctuating force in Brownian motion is “white”, i.e. the fluctuations at different instants of time are not correlated and the corresponding power spectrum is flat over the whole range of frequencies. This alone (and the value of the classical friction constant) fixes the important low-frequency behaviour of the bath spectrum also for the quantum case, at low temperatures.

This type of bath, leading to a velocity-proportional friction and a white-noise force at higher temperatures is called “Ohmic”. That is because the Nyquist noise of the electric field (and the current) in an Ohmic resistor is also “white” at low frequencies and finite temperatures.

(Offline readers: There is a Java simulation of a classical bath of oscillators available at iff.physik.unibas.ch/~florian/cl/dissipative.html).

Dephasing vs. dissipation: The example of the harmonic oscillator

If one applies the model of Quantum Brownian motion to the harmonic oscillator, the result may look like the one shown in the following picture.



In this picture, the probability density is plotted in a space-time diagram, where time runs to the right and the oscillator coordinate is plotted on the vertical axis. The oscillator starts in a superposition of two Gaussian wave packets, which are at two opposite sides of the origin. In the course of time, the center of mass positions of the packets oscillate just like a classical oscillator would do. Whenever they meet at the origin, an interference pattern becomes visible. It is due to the superposition of the plane

wave components which make up the two packets. The momenta of these plane waves are centered around the classical value of the momentum of an oscillator swinging through the origin. Since both oscillations move in opposite directions, the momenta are $+p$ and $-p$, so the interference term in the wavefunction has a wavenumber of $2p/\hbar$, giving rise to the fringes observed in the picture.

On the time-scale shown in this picture, the *amplitude* of the oscillations stays constant approximately, so the friction force has not yet led to dissipation of an appreciable fraction of the oscillator energy. However, the interference patterns clearly become “washed out” in the course of time. This is due to the dephasing effects of the “Ohmic bath”. Obviously, dephasing proceeds much faster than energy dissipation in this example. This can be understood quite easily: In fact, if a quantum oscillator is in excited level number n and makes a transition to lower level $n - 1$ (due to the action of the bath), the phase of the quantum state is already lost completely. However, in terms of energy, this corresponds only to a small fractional change, namely $1/n$. Therefore, the dephasing rate is n times as large as the energy decay rate. In a very rough but qualitatively correct fashion, this also explains why dephasing is so fast for superpositions of “macroscopic” or “nearly classical” quantum states: These belong to the limit of large quantum numbers.

Quantum chaos: Adding noise helps

The model of “Quantum Brownian motion” has also turned out to be useful in the field of quantum chaos. There one wants to find out what happens to classically chaotic motion when quantum mechanics becomes important. The answer generally is that quantum mechanics opposes chaos.

However, in real life there is some “noise”, i.e. fluctuations which can dephase the quantum-mechanical coherent motion. It turns out that they are able to restore the features of the classically chaotic motion, so that, for example, initial states which are quite close to each other get separated exponentially fast in the course of time. Amazingly, the exact amount of noise may be unimportant, if it lies within some limits! It can be strong enough to restore the classical features of the motion, yet weak enough so

that the friction and heating which invariably go with the noise do not matter (on the time-scales considered). In order to describe these effects of dephasing on quantum chaos, one often uses the model of “Quantum Brownian motion”, since it is somehow “generic” and the calculations get comparatively simple in the limit of high temperatures.

Quantum Brownian motion at low temperatures

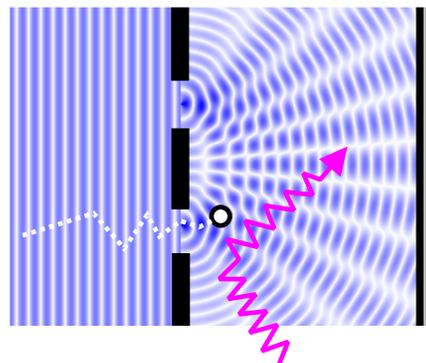
Quantum Brownian motion becomes particularly interesting at low temperatures. The kind of bath spectrum which one has to choose in order to get a white noise force at high temperatures produces a force with long-range correlations in time at low temperatures. Therefore, the situation is quite different from classical Brownian motion, where the force acting on the particle at one instant of time is independent of the history of the particle’s motion. Although the particle still relaxes to a well-defined equilibrium state independent of its initial conditions (“ergodic motion”), it does so “much more slowly” than at higher temperatures. In fact, at $T = 0$, the spread of an ensemble of particles undergoing quantum Brownian motion only grows logarithmically with time - compared with the normal diffusion at high temperatures, where the width grows like the square root of time. There are also peculiar properties of *dephasing* in this case. For example, the loss of some interference pattern which had been present in the beginning is not described by the usual exponential decay law with some “dephasing rate”, but instead it is “only” a power-law decay in time. Compare this to the case of the atom, where the excited level has a finite decay rate at $T = 0$. The difference stems from the fact that the quantum levels of the free particle are spaced infinitely close and the decay rate depends on the level. Thus, when the relaxation proceeds from highly excited states towards lower states, the corresponding decay rates get smaller and smaller, resulting finally in the power-law decay of the interference pattern, instead of an exponential decay.

It is not entirely self-evident that coupling to the bath leads to relaxation to equilibrium at $T = 0$. For example, a freely moving charge is not affected at all by the electromagnetic field and its vacuum fluctuations, due to Lorentz invariance: Any inertial frame of reference is

equally good, so there is no reason for the particle to slow down. Put differently, only an accelerated charge (like the electron in the atom) will radiate away energy. This example shows that there are baths which do not always lead to ergodic behaviour at zero temperature. On the other hand, for Brownian motion it is to be expected, due to the velocity-proportional friction force that decelerates the particle.

The “Heisenberg microscope”: Finite wavelengths of the fluctuating field

In our discussion up to now, the fluctuating force has had the same magnitude and direction throughout all of space. Its fluctuations only take place in time. One may wonder what is the effect of a force that is not constant in space, i.e. where the wavelengths of the field are not assumed to be very large. That this is indeed an important question may be seen in the example of the “Heisenberg microscope”:



Heisenberg’s gedanken experiment is about finding out the position of an electron - for example, to determine which slit the electron goes through in the double-slit experiment. One could try to do this by scattering light off the electron just when it is about to go through the slits. Of course, a high spatial resolution is needed to determine exactly through which of the slits it has gone. In order

to achieve this resolution, one must use light of a very small wavelength, smaller than the distance of the slits. However, although it could (in principle) be done and although the electron will always be detected in either one or the other slit (not both at the same time), it does not help to clear up the mystery of how the interference pattern is actually formed. This is because the interference pattern vanishes whenever the spatial resolution is large enough to detect which way the electron goes through the double slit setup. The photons of small wavelengths have large momenta, so they will change the momentum of the electron drastically in the course of scattering. Since the momentum transferred to the electron is random, the deflection of the electron is random as well. Therefore, a lot of different interference patterns are superimposed on the screen, each of them belonging to a different momentum transfer. Taken together, they just form a broad distribution without any interference fringes. The interference pattern gets completely smeared out just at the point when the wavelength of the photons gets small enough to detect which slit the electron has gone through. This is a consequence of Heisenberg's uncertainty relation, which this gedanken experiment was invented to illustrate.

What do we learn from this about dephasing? If an electron in an interference setup absorbs (or emits) a field quantum of a sufficiently small wavelength such that its path could (in principle) be resolved, this is enough to destroy the interference. Conversely, it may happen that field fluctuations do not dephase the particle's motion when their wavelength is too large - even if they are rather strong!

Perhaps it helps to look at an example related to the Heisenberg gedanken experiment: Suppose we do not shine light on the electron to find out its position. Will it not, nevertheless, reveal its position automatically, by radiating away photons? After all, the vacuum fluctuations of the electromagnetic field should not only induce an atom to decay from its excited state but should also be able to dephase the electron in this setup. As we have pointed out above, a freely moving electron will not be harmed by those vacuum fluctuations and it will not emit photons, simply because it is not accelerated and its inertial frame is as good as any. But at the slits (or any mirrors etc. in the setup) it will have to change direction (with a certain probability at least) and this scattering by a poten-

tial should result in "braking radiation" (which also produces the X rays in an X ray tube). In fact, the emission of photons will even happen with unit probability! However, most of them are very long-wavelength photons. (In earlier times, this fact was called the "infrared catastrophe") They never reveal through which slit the electron has gone, only perhaps that it is present in the interference setup at all. The correct question to ask is: what is the probability for the emission of photons that have a wavelength short enough to distinguish between the two interfering paths? The answer is that, under reasonable experimental conditions (electrons moving with less than 10% of the speed of light), this probability is less than one in ten thousand, so the interference pattern will not be spoiled at all. Furthermore, thermal fluctuations of the electromagnetic field are also completely irrelevant, because, again, their wavelengths are much too long (if one does not perform interference experiments inside a very hot furnace...).

Measurement means dephasing and entanglement

In the previous example of the Heisenberg microscope, we have already pointed out the fact that the measurement of the electron's position destroys the interference pattern, i.e. leads to dephasing. This is generally true: As soon as the different interfering paths become distinguishable in principle, the interference is lost. It is the coupling to a measurement apparatus or a bath which makes the paths distinguishable. In fact, for this purpose, a bath is equivalent to a measurement apparatus and vice versa. The difference is really a practical one: While the apparatus is designed such that the information acquired about the system can be read out reliably by the experimenter, it is simply lost in the case of interaction with a bath.

In the quantum mechanical formalism, the possibility of distinction between the paths of the electron, or generally the states of the system, is due to **entanglement** between the quantum states of system and bath (or apparatus). Because this is so central to dephasing, we will explain it now briefly (even though the aim of this introduction is to avoid the formalism and emphasize the physics).

One starts out with a system in a quantum state $|\Psi\rangle$ which is a superposition of two states $|1\rangle$ and $|2\rangle$:

$$|\Psi\rangle = \psi_1 |1\rangle + \psi_2 |2\rangle$$

Here the complex numbers ψ_1 and ψ_2 are the amplitudes of the two states, whose squares add up to give the total probability one: $|\psi_1|^2 + |\psi_2|^2 = 1$. The two states could be “spin up” and “spin down”, “atom excited” or “deexcited” or an electron “going the upper way” and “going the lower way” (roughly speaking).

If a “bath” (or an apparatus) is present, the total wavefunction is that of *both* system *and* bath. If they have not yet interacted, their states will be completely independent. For example, the bath may be in its ground state (the “vacuum” in the case of the electromagnetic field), independent of the system’s state. In this case, the total wavefunction is just the product \otimes of the system state shown above and the bath state $|\chi_0\rangle$:

$$(\psi_1 |1\rangle + \psi_2 |2\rangle) \otimes |\chi_0\rangle$$

After the interaction has taken place for a while, the total quantum state is no longer simply a product state - and therefore it is called “entangled”:

$$\psi_1 |1\rangle \otimes |\chi_1\rangle + \psi_2 |2\rangle \otimes |\chi_2\rangle$$

This is the simplest example of an entangled state that may result in the course of the interaction. The system state has not changed but the bath state is changed: It is no longer the ground state $|\chi_0\rangle$, but has become a different state, *dependent* on the system’s state - either $|\chi_1\rangle$, if the system had been in state $|1\rangle$, or $|\chi_2\rangle$, if the system had been in state $|2\rangle$. Thus, the bath state now carries some information on the system’s state. In the case of the measurement apparatus, one speaks of a “pointer” state, because it acts like a needle in a conventional instrument. Since the system started out in a superposition of states, the total wavefunction is now also a superposition of the two different possibilities.

How does this destroy an interference pattern? First one must ask how “interference” shows up in the initial superposition of system states. Generally, interference means

dependence of some observable quantity on a phase difference. For example, suppose states $|1\rangle$ and $|2\rangle$ are two orbitals of an atom (or plane wave states of an electron). Then, the probability density of finding the electron at a certain point x in space not only depends on the densities of the individual orbitals, but also on an interference term. As usual, the probability density is given by the magnitude squared of the projection of the electron state $|\Psi\rangle$ onto a position state $|x\rangle$:

$$\begin{aligned} |\langle x | \Psi \rangle|^2 &= |\psi_1|^2 |\langle x | 1 \rangle|^2 + |\psi_2|^2 |\langle x | 2 \rangle|^2 \\ &\quad + 2 \operatorname{Re} [\psi_1^* \psi_2 \langle 1 | x \rangle \langle x | 2 \rangle] \end{aligned}$$

Now if the complex amplitude of, say, the second state, ψ_2 , is multiplied with a phase factor $e^{i\varphi}$, this affects the value of the interference term, i.e. the term which contains both complex amplitudes ψ_2 and ψ_1 . The total probability density thus depends on the phase-difference between these amplitudes.

After the interaction with the bath, the interference term is replaced by something containing also the overlap $\langle \chi_1 | \chi_2 \rangle$ of the two bath states:

$$2 \operatorname{Re} [\psi_1^* \psi_2 \langle 1 | x \rangle \langle x | 2 \rangle \langle \chi_1 | \chi_2 \rangle]$$

This additional complex number may, first of all, change the phase of the interference term. But more important is the fact that its magnitude can be shown to be always smaller or equal to unity. If it is smaller, the interference term will be suppressed relative to the other terms. In the extreme case where the overlap vanishes, there is no interference any more: Any dependence on the phase difference between the complex amplitudes ψ_1 and ψ_2 is lost completely. That takes place exactly when the information acquired by the apparatus or bath is, in principle, detailed enough to distinguish perfectly between the two system states (just by determining the bath state itself).

You can imagine that in general things may get more complicated: The system’s state may change as well in the course of the interaction. For example, in the case of the atom, this may go so far that in the final superposition both states $|1\rangle$ and $|2\rangle$ are replaced by the ground state, so

the total wavefunction is a product state again and not entangled anymore. Then, the quantum state has been transferred entirely from the atom to the field, which is now in a superposition of “1 photon there” and “no photon there”. In addition, the two parts of the final total quantum state need not be simple product states themselves (which they have been in the example): An atom in an excited state may go either directly to the ground state or to another lower-lying level, each with a certain probability. However, the general principle is just as shown here.

Entanglement does not automatically mean dephasing!

There is one possible misunderstanding which has to be avoided in this context: In general, it is *not* true that entanglement between system and bath alone leads to dephasing. There are important situations when a small system is entangled with a bath to an appreciable degree, and yet shows perfectly coherent oscillations or interference patterns.

This will occur whenever the bath is able to adapt without delay to the system’s state, i.e. when it follows the motion of the system adiabatically (because it consists of modes with high frequencies). Take, for example, the coupling of an electron to the vacuum fluctuations of the electromagnetic field. When the electron goes through some interference setup, the total state is not just a superposition of “electron here or there”. Rather, the configuration of the electromagnetic field is changed in the vicinity of the electron. This means that the state of the electromagnetic field depends on the position of the electron, such that the total state is indeed an entangled state! The same happens for the motion of an electron in a crystal, where it will distort the surrounding crystal lattice. However, this has almost no effect on any interference pattern. Whenever the two possible paths of an electron meet at some screen, where the pattern is to be observed, the corresponding states of the field coincide again, so their overlap is unity. In fact, there *is* a slight effect of the coupling to the bath: The “vacuum fluctuations” of the field lead to some jitter of the electronic position. This is, after all, the origin of the Lamb energy shift in an atom. However, the effect is usually small and, more importantly, it does not grow with increasing path length travelled by the electron - quite in

contrast to “real” dephasing, where coherence properties are destroyed more and more in the course of time.

In the case of a two-level system, coherent oscillations between system states $|1\rangle$ and $|2\rangle$ will just be replaced by coherent oscillations between entangled states $|1\rangle \otimes |\chi_1\rangle$ and $|2\rangle \otimes |\chi_2\rangle$. Although the coupling to the bath usually changes the oscillation frequency and also the magnitudes of some quantities (like the average dipole moment in the case of the atom), it does not destroy the oscillations themselves.

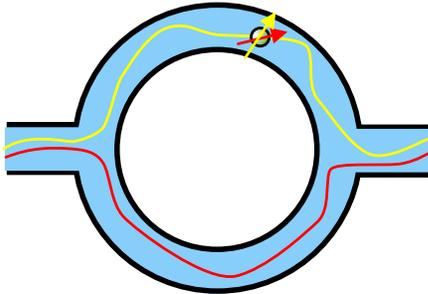
In principle, this is nothing but the story of “bath fluctuations which are too fast” to dephase the system, only dressed in the language of entanglement (and applied to zero temperature).

Dephasing will take place only due to those modes which have frequencies comparable to the system’s motion. At zero temperature, the system will just emit quanta into these modes, which carry away both energy and information about the system state. Since they are free to move away and be detected somewhere else, they are called “real” quanta, in contrast to the “virtual” quanta which belong to the field state that follows adiabatically the motion of the particle. The “virtual” quanta are tightly bound to the particle and cannot be set free due to energy conservation. They are only present as long as the interaction is there, while “real” quanta live on even when the interaction that produced them ceases again. In this sense, the type of entanglement discussed here is “reversible”. Apart from that, there is (of course) no difference between both types of quanta.

The Aharonov-Bohm ring: “which-way” information, energy transfer and the Pauli principle

A simple version of the double-slit experiment for electrons consists in a metal ring where an electron can travel either the upper or the lower arm (see the picture). This has the advantage that the phase-difference between the two paths can be adjusted by applying a magnetic field which threads the ring. Even when the field is confined to the interior of the opening in the ring and cannot affect the electron’s motion directly, the Aharonov-Bohm effect will make the phase-difference dependent on the magnetic flux

inside the ring. In the experiments, the diameter of such a ring is typically only a few micrometers.



An Aharonov-Bohm ring, with the two possible paths of an electron that goes through the ring once. According to the path it takes, the spin may end up in one or the other final state (indicated by colors).

Suppose some two-level system (for example: spin or tunneling system) is present in one of the arms of the ring. It will be influenced if the electron moves through that particular arm. The spin may be flipped or the particle may jump from the left to the right well of the double-well potential. In this way, the electron has left a trace in its “environment”, revealing the path it took. In that case, every interference effect that depends on the phase-difference between the two arms is lost, since the electron has become entangled with the environment and the new state of the latter is orthogonal to the old one. Of course, this only happens with some probability, depending on the strength, duration and type of the interaction between the electron and the two-level system. To put it differently: The new state need not be fully orthogonal to the old one (the spin has been rotated only a little bit), so the interference pattern is not lost completely, as has been explained above.

Since the states of the two-level system will in general have different energies, the electron may need to supply enough energy to flip the state. This would be no problem at all if the electron were the only electron around and moving with some rather high velocity. Then, we would have practically the same situation as in Heisenberg’s gedanken experiment, with a single electron passing through a double-slit setup. The only difference would be that the electron’s path is revealed automatically - it triggers some measurement when going through one of the slits. However, in a metal the electron is not alone.

This means that it cannot just emit an arbitrary amount of energy into the “bath”, because it may end up in a state which is already occupied by another electron and this is forbidden by the Pauli principle.

At temperatures that are much higher than the energy needed to flip the spin, this is not really a problem. Each of those fast-moving electrons that carry the electric current can then easily emit this much energy, because the chance that it will end up in an occupied state is very low: The states are only sparsely occupied in a “thermal range” around the Fermi energy. On the other hand, at such temperatures it may not even be necessary for the electron to *supply* the energy. The two-level system will be in its excited state with a 50-50 chance and then it will pass energy to the electron when being flipped to its ground state.

At low temperatures, when the electron cannot flip the state of the spin or tunneling system, the two still become entangled - but this entanglement only persists as long as they interact. This is the kind of “reversible” entanglement which has been discussed above. After the electron has moved on, the two-level system will be left in its original state and therefore this does not lead to real dephasing.

The “Pauli principle mechanism” indeed limits the strength of dephasing which an electron experiences in a solid at low temperatures. Only if the two-level system’s states have degenerate energies, it will be able to dephase the electronic motion all the way down to zero temperature because then no energy transfer is needed to flip the state. Of course, the same qualitative considerations apply to any other type of environment, be it phonons or even the other electrons. Then, the possible energy transfers are distributed according to the bath spectrum.

Whether dephasing is really suppressed by the Pauli principle does not only depend on the temperature but also on the type of experiment carried out. If, for example, one drives an electric current with a rather *high* voltage, this may mean that the electron is supplied with enough energy to flip the spin, leave a trace in the environment and lose its phase - i.e. to destroy the interference effects. But in a measurement of the resistance (or its inverse, the conductance), the voltage applied is ideally infinitely small, so that possibility is ruled out in such a situation.

In any case, at low temperatures the Pauli principle plays an important role in the dephasing of electrons in solids

and thus adds an additional feature which is completely absent from all *single*-particle interference experiments (whether in optics or with single atoms, ions or electrons). This makes the discussion rather subtle, because one is not allowed any more to think in terms of a single particle that is merely coupled to a bath. In fact, since every electron couples to the bath, this will in general even lead to an effective interaction between the electrons themselves, so that, apart from the original bath, one has to consider the other electrons as an additional environment - even if one starts out with a model where the original Coulomb interaction between the electrons had been neglected.

Interference effects in mesoscopic physics and optics: the role of dirt

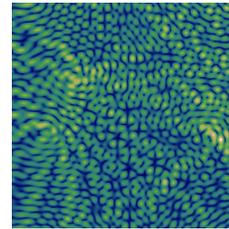
The example of the Aharonov-Bohm ring has already demonstrated that in solid-state physics - or more precisely mesoscopic physics - interesting effects may become important in discussions of interference and dephasing.

In mesoscopic physics, one analyzes electronic circuits on the micrometer scale at low temperatures, where the wave properties of the electrons become important. Interference effects show up everywhere: The flow of electrons is now described much better in terms of the analogy to the passage of electromagnetic waves in a network of waveguides than using Ohm's law. For example, the current running through one wire of the circuit depends on the presence of other leads attached to it even when these do not carry any current themselves.

Sometimes, the interference effects are quite "fragile". For example, the electric current is carried by the fast electrons at the Fermi surface, so it is *their* wavelength that determines the spacing of interference fringes. However, at finite temperatures, the range of relevant energies (and wavelengths) gets smeared around the Fermi energy, so that one observes an average over interference patterns belonging to different wavelengths, which appears washed out. Dirt is another problem: A single impurity atom that scatters the electron waves may completely shift the phase of the electron going through a wire. If the experimental setup is such that only the *average* over the current measured in many "identical" systems is observed (which dif-

fer in the detailed placement of the impurity atoms!) then again the interference pattern may disappear.

However, it would not be correct to think that finite temperatures or dirt alone destroy *all* the interference effects. In principle, impurities only produce a very complicated but still perfectly coherent interference pattern, which may look like the one in the picture below:



This is completely analogous to the scattering of coherent monochromatic light from a disordered material: There, the reflected light shows a "speckle pattern" which is due to the complicated superposition of all the spherical waves emanating from the randomly placed scatterers.

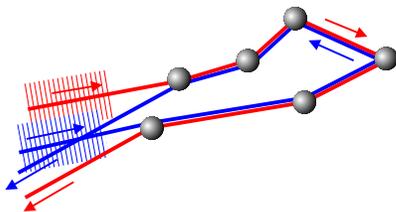
Even moving a *single* impurity can produce an important change in the interference pattern and, therefore, in quantities like the conductance of a sample. The same will happen if the impurities remain fixed but the wavelength of the electrons is changed slightly or a magnetic field is turned on that introduces extra phases for the wave motion of the electrons. Experimentally, the last two possibilities are easier to realize, of course. (The wavelength can be tuned by changing the density of the electrons) If one varies such an external parameter like wavelength or magnetic field, the conductance will fluctuate wildly, depending on the detailed arrangement of impurities. Interestingly, the *magnitude* of the fluctuations in the conductance is more or less always the same: This is why this phenomenon is called "**universal conductance fluctuations**".

If one takes the average of the conductance over many samples (with different placement of impurities, but otherwise identical), nothing remains of the conductance fluctuations or, in general, any interference effect that had been sensitive to the details of the impurity configuration. Observing things at finite temperatures is similar, because it means averaging over different wavelengths.

Nevertheless, there are interference effects which still

show up even when only an average over many disordered samples is considered.

Such an effect was first observed in the scattering of radar waves from clouds, where the intensity of the diffusely reflected radiation was found to be enhanced in the backward direction. The droplets in the clouds act like randomly placed scatterers and the same effect is observed in the scattering of *any* kind of radiation from a disordered medium: light, sound waves, electron waves etc.! However, the strength of the effect depends on the relation between the wavelength λ and the mean free scattering path l , which describes how far the wave can travel on average before being scattered. If scattering is weak, so that l is much larger than λ , the effect is tiny (but sometimes still important). Since the effect depends on the interference of waves that have been scattered from several scatterers (“multiple scattering”), it only occurs if the phase is not scrambled on the way. The radiation should remain coherent during the time it takes to make a round-trip visiting several scatterers. That is why the effect is called “**coherent backscattering**”. For the electrons, it means that the conductance is lowered, since the backscattering is enhanced. This is called “**weak localization**”, since it is the precursor of complete (“strong”) localization of electron waves in a disordered medium, when the waves cannot move at all and the electric conductance vanishes. (Strong localization occurs predominantly in one- or two-dimensional systems)



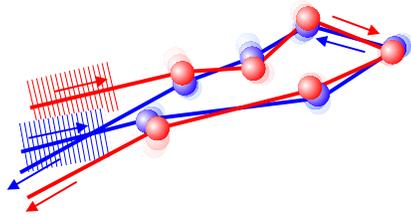
Multiple scattering of a wave from impurities, leading to weak localization: One of the paths is the time-reverse of the other.

Calculating the scattered intensity is in principle straightforward: One “simply” has to sum the amplitudes for scattering at any number of scatterers, in every possible order. The magnitude squared of the resulting amplitude gives the intensity. Expanding out the square of the sum over paths obviously gives a sum over *pairs* of paths (i.e. a sum over products of two amplitudes each). This is

completely analogous to the treatment of interference in a double-slit experiment, where one has to consider the phase difference between the two possible paths. However, now there is not only one pair of paths but one has to sum over the contributions of every possible pair of those “random-walk” paths. Usually, the phase difference fluctuates wildly from one pair to the next, so the contributions average out. There are two possible special types of pairs where the phase-difference is always zero, so this argument does not apply: The first one corresponds to pairing the path with itself and gives a contribution as one would expect it from a classical calculation, where one considers *probabilities* of scattering instead of amplitudes. The other one corresponds to pairing a path with its time-reversed copy, where the scattering takes place at the same scatterers, only in the reverse order. There again, the total phase accumulated along the two versions of the path is the same. This contribution gives the extra “coherent backscattering”. It is represented in the picture above. Although it is of the same size as the classical contribution, it reacts much more sensitively to external disturbances (dephasing) or a magnetic field. In these situations the phases of the path and its time-reversed copy will become different.

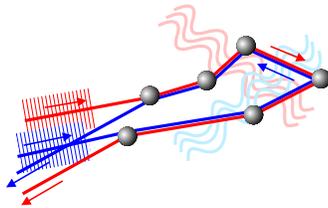
Dephasing of “universal conductance fluctuations” and in “weak localization”

The interference effects described above can be destroyed in a variety of ways: For coherent backscattering of radiation from atoms in an atomic vapour, it suffices if the atoms are moving about, such that the wave will encounter another configuration at later times. Then the phases accumulated along a path and its time-reversed copy are different, and the difference itself fluctuates (see the picture). The longer interference paths obviously get their phase scrambled more easily, since it takes more time to traverse them. This then results both in a suppression of the intensity of backscattered radiation as well as a “smoothing” of its angular dependence (in the ideal case the intensity of backscattered light shows a sharp cusp in the backward direction when plotted versus the scattering angle). If dephasing is very strong, the effect is destroyed completely, such that only the diffuse “single” scattering remains, which is usually isotropic.



Dephasing due to motion of the scatterers: The partial waves traveling along one path or its time-reverse see a different configuration of scatterers.

For electrons in a solid, the impurities themselves usually do not move. However, there may be *magnetic* impurities, i.e. spins to be flipped and “revealing” the path of the electron. Then, there are the vibrations of the crystal lattice which constitute a “bath” that dephases the electron’s motion especially at higher temperatures (shown in the picture: the different distortion patterns of the lattice, acting on the path and its time-reversed copy). At low temperatures, the “Nyquist” fluctuations of the electric field due to the random motion of other electrons become more important. This is simply a consequence of the fact that the bath spectrum of the electron bath is stronger at low frequencies than that of the phonon bath and that, at low temperatures, only the low-frequency bath modes matter. After all, the higher-frequency modes are neither thermally excited nor has the electron enough energy available to excite them, due to the Pauli principle.



Schematic representation of dephasing by electrons or phonons in a metal: The electron wave feels a different fluctuating electric field depending on whether it travels one or the other path.

Because of dephasing, the conductance fluctuations get washed out, so that their magnitude is diminished and the fine structure in the plot of conductance versus magnetic field or wavelength is no longer visible.

The weak localization effect is also reduced, which means the resistance is *lowered*. Usually it would not be possi-

ble to measure this rather tiny resistance change. However, measuring the resistance as a function of the magnetic field (“magnetoresistance”) is one accurate way to find out about the strength of dephasing. This is because the magnetic field itself destroys the weak localization effect. The two time-reversed paths acquire different phases in a magnetic field, which depend on the magnetic flux enclosed by the path (Aharonov-Bohm effect). Since the paths are random, this phase difference fluctuates wildly and the sum over all paths does not show any more the weak localization interference effect. Again, the effect of the magnetic field is more pronounced for longer paths, since they enclose a larger magnetic flux. If the magnetic field gets so weak that it could only affect very long paths which are dephased anyway (by phonons, electrons, magnetic impurities etc.), it ceases to have an influence on the resistance. This is how the strength of dephasing can be obtained in experiments - for example, to analyze it as a function of temperature.

Dephasing by Nyquist noise: The numbers

It is instructive to look at the typical numbers of the dephasing time in this example and to understand how they come about. At low temperatures, the most important effect is due to the Nyquist noise of the electric field, which is produced by the random motion of electrons in the metal (- piling up a little bit extra negative charge here or there by chance). If one looks at these fluctuations with a rough resolution, the values of the electric field at different points in space and time fluctuate completely independently (“white noise”). The average square of these fluctuations is proportional to the temperature and to the resistivity (which is plausible). If you ask for the precise magnitude, however, you have to specify over which volume in space and over which time interval you want to average the fluctuating field (specifying the resolution). As usual, the typical magnitude of the *averaged* fluctuating field drops with the inverse square root of the volume and the time-interval, just as the statistical error of an average drops with the inverse square root of the number of data points considered. Which volume and which time-interval should one insert here to get an estimate of the field fluctuations that will dephase the electronic motion? Consider an electron moving along one of those random

paths. If the path extends at most about a distance L away from its starting point, then L^3 is the volume over which the average of the electric field fluctuations is to be taken. Of course, in a thin wire of cross-section A , the volume is $A \times L$ instead, because the electron cannot move outside of the wire in the transverse direction (perpendicular to the wire axis). The time is simply given by the time it takes to travel through the path. Since this is a diffusive motion, traveling double the distance quadruples the time needed.

The average magnitude of the electric field fluctuations due to Nyquist noise turns out to be somewhere around one V/m at a temperature of one Kelvin, when one plugs in the typical numbers of one micrometer cubed volume and a nanosecond time interval (as well as typical resistivities). In contrast to that, the vacuum electric field fluctuations which give rise to spontaneous decay of an atom in free space have a strength of roughly $10^4 V/m$. Therefore, the Nyquist noise may seem comparatively tiny. However, remember that the decay of the atom is governed by vacuum fluctuations which are at frequencies around $10^{15} Hz$. Here, we are dealing with typical times of a nanosecond, corresponding to a million times lower frequencies. At such a low frequency, the vacuum (free-space) electric field fluctuations are already down by a factor of 10^9 , much too small to play any significant role in the dephasing of electrons in metals.

While the fluctuations of the averaged field get smaller if one increases the path length, the resulting fluctuations of the phase difference get larger. Calculating the phase-difference between the two different paths considered (a path and its time-reverse) means integrating the difference of energies over time. The difference of energies is the difference of electrostatic potentials at two points (on the two paths). This potential difference is the value of the electric field multiplied by the distance, which is typically L again. The travel time itself goes like L^2 , as we have already seen above. If you have counted all powers correctly, you will find that the average *square* of the phase-difference grows like $L^2 \times (L^2)^2 \times T / (L^3 \times L^2) = T \times L$ in the three-dimensional case and like $L^2 \times (L^2)^2 \times T / (A \times L \times L^2) = T \times L^3 / A$ in the case of the wire. As soon as the fluctuations of the phase-difference reach one, the phase coherence of

the paths is essentially lost. This then defines the phase coherence length L in the example of weak localization - and prescribes its theoretical dependence on temperature. Specifically, in the important case of the wire, the length L goes like $T^{-1/3}$, which means the dephasing time, being proportional to L^2 , grows as $T^{-2/3}$ when one lowers the temperature. (In the three-dimensional case it turns out that high-frequency fluctuations of the electric field, which have nothing to do with Nyquist noise, are the dominant ones - and their temperature dependence is different from the one derived here.)

Using the typical values of resistivity and diffusion constant, one arrives at a phase coherence length in the micrometer range at a temperature of about one Kelvin (that's why we used these values already above). This is also what experiments find. The time it takes the electron to travel random-walk paths of this length is about one nanosecond. Raising the temperature decreases the dephasing time, so that it will be down by a factor of about 20 at 100 Kelvin, meaning that the phase coherence length drops to about one fourth its value at one Kelvin - if it were not for the crystal vibrations, which accelerate the loss of coherence at higher temperatures.

A dephasing time of one nanosecond means that the motion is impressively coherent: After all, during that time interval the electron will have scattered about ten thousand times from impurities, so the "multiple scattering" at the heart of the weak localization effect is fully developed. The micrometer coherence length is about one hundred times a mean free scattering path length.

Puzzling behaviour at low temperatures

When the temperature becomes low, the fluctuations in the environment of an electron get less severe. Therefore, dephasing is weakened and one starts to observe all the different interference effects described above - those that depend on the specific arrangement of impurities and the geometry of the sample, as well as those that are more robust in that they remain after averaging over impurity positions. The electrons remain phase-coherent over longer and longer periods of time, in which they can travel farther, so the interference effects show up in larger and larger parts of the sample. In fact (as it was explained

above), the theory of weak localization predicts an increase of the dephasing time beyond all bounds as one approaches zero temperature - at least as long as strong localization does not set in and there are no magnetic impurities.

Therefore it is no surprise that the demands on experimental accuracy become greater as one lowers the temperature. If the experiment is not carried out very carefully, some external noise (perhaps due to the measurement setup or some not-quite-static control fields) may lead to extra dephasing which then dominates any “intrinsic” dephasing mechanisms that have already become quite weak. It also becomes more important to approach the ideal limit of zero voltage in measurements of the linear conductance, otherwise the electrons will be supplied with “surplus” energy that they can pass into a bath, thereby destroying interference effects. Joule heating and shot noise may be other problems when one carries out the experiment at a finite current. In weak localization measurements, any stray magnetic fields or remaining magnetic impurities in the sample may as well spoil the experiment. And finally, since one is interested primarily in the temperature dependence of the “dephasing time” (observing its increase towards lower temperatures), one has to take care to measure the temperature itself accurately, which is not so easy at low temperatures, when different parts of the system (electrons, phonons) may be at different temperatures.

For all of these reasons, it had not been considered a serious trouble that experiments in mesoscopic physics consistently showed a saturation of the dephasing time at low temperatures, in contrast to the predictions of theory for the ideal case. However, some years ago new weak localization experiments were carried out in which special care was taken to eliminate the known culprits for an “extrinsic” saturation of the dephasing time. It was found that the saturation at less than about one Kelvin remained, so the dephasing time did not get larger and larger as “demanded” by theory. It was speculated that zero-point fluctuations of the electric field in the metal (not the free-space variety) might be responsible for the dephasing time leveling off at low temperatures. Consequently, the theory has been reexamined, paying special attention to the fact that the arguments on energy conservation and the Pauli principle given above can be justified directly only

in a perturbative calculation. However, at present, the interpretation of the experiments is still unclear. Even if the saturation turns out to be “only” due to some remaining external noise, the experiments would show practical limitations to the observation (and exploitation) of single-particle interference effects in mesoscopic physics. Furthermore, it has been suggested that the same noise could be at the origin of the surprisingly large size of the “persistent” currents which are observed to flow at equilibrium in small metal rings at low temperatures.

Summing it up

What can one learn from the concept of dephasing (or “decoherence”)? First of all, dephasing sets the limits to the observation of interference effects, whether they are coherent oscillations of simple two-level systems (“qubits”), coherent dynamics of ions in traps, freezing of chaotic motion in quantum chaos, “classical” interference setups or the subtle phenomena encountered in the scattering of waves in disordered media. It also provides a theoretical model problem which is midway between single-particle physics and the full-fledged many-particle problem. Observation and analysis of the various “dephasing times” can reveal details about the “bath” and the interaction between particle and bath. And on a more general level, dephasing helps to explain why the weird effects of quantum mechanics are not observed in the behaviour of macroscopic bodies.



Further reading

The book by Weiss, “*Quantum-dissipative systems*” (World Scientific, 2nd edition 1999), gives a thorough introduction to the theoretical treatment of dissipation and decoherence in quantum mechanics (with particular emphasis on the two-level systems). I have based the discussion of an atom’s spontaneous decay on the very readable account in Milonni’s book “*The Quantum Vacuum - an introduction to Quantum electrodynamics*” (Academic Press, 1994). Quantum diffusion experiments are reviewed in the article “*Coherence and Decoherence of*

Positive Particle States in Solids”, by Karlsson, in *Physica Scripta* T76, 179 (1998). The basics of mesoscopic physics (including dephasing in weak localization, the example of the Aharonov-Bohm-ring and the relation to entanglement) can be found in Imry’s book, “*Introduction to Mesoscopic Physics*” (Oxford University Press, 1997). Another reference with particular emphasis on dissipation and dephasing in mesoscopic systems (including quantum chaos) is: Dittrich et al., “*Quantum Transport and Dissipation*” (Wiley-VCH, 1998). References to the recent debate on saturation of the weak-localization dephasing time at low temperatures can be found, for example, by using WebOfScience to locate citations to the original letter describing the measurements (Mohanty, Jariwala and Webb, PRL 78, 3366, 1997).