The first TOPOLOGICAL quantum programming language

A conversation with Urs Schreiber

The Mansions of Science · Original video

Edited from the supplied automatic transcript. Fillers, stutters, false starts, and redundant repetitions have been removed; sentences and paragraphs have been reconstructed for readability. The interview's sequence, substantive exchanges, examples, and qualifications have been retained. Speaker attribution is inferred from the text. Clear caption errors have been corrected, and Urs Schreiber has clarified the seven passages initially flagged for review, supplying names, titles, and references where needed. This edition has not been checked against the recording. Chapter times follow the supplied video description.

1. Familiarity and understanding · 0:00

Interviewer: Thank you so much for your time. Before we get started, I want to read something you said to *The Reasoner* a while back:

> *One striking aspect of higher topos theory is that it makes deep phenomena conceptually simpler. Often people regard higher topos theory as hideously complicated or scary. But this is a deception, conflating simplicity with familiarity. It is similar to how people on the street will say it is simple to understand why apples fall to the ground, but hard to understand what happens in a nuclear power plant. What they really mean is that they are used to apples falling to the ground. In the actual sense of understanding, the theory of gravity is still mysterious, while the nuclear forces have been understood to great depth.*

I want anyone watching or listening to keep that distinction in mind: familiarity and understanding are different things. These supposedly scary toposes can help with understanding, even though they are unfamiliar.

2. Higher gauge theory and current research · 1:28

Interviewer: What are you working on right now, and what are you excited about?

Urs Schreiber: We have an article coming up, hopefully next week, on deriving K-theory from M-theory. I am quite excited about that. We also have projects more closely related to condensed matter theory.

Broadly, my work concerns nonperturbative aspects of quantum field theories: the global topological structures of, in particular, higher gauge theories, such as those that appear in higher-dimensional supergravity. We apply these ideas to mathematical physics generally, but in recent years we have also been making applications to actual experimental physics.

Interviewer: Let me explain my understanding of higher gauge theory. In ordinary gauge theory, suppose we have a manifold-say, four-dimensional spacetime. At every point we attach a fibre, such as a circle, S^1, corresponding to the group U(1). As an electron moves through the manifold, from one fibre to the next, its phase changes.

In your papers, you replace groups with 2-groups, and bundles become 2-bundles. Is that correct? And why do you do it?

Urs Schreiber: That is roughly correct, but it needs some commentary. Let me start with why.

When I began looking into higher categorical structures, I had no background in them. I had a PhD in theoretical physics and was working on string theory. String theory famously contains what are sometimes called higher tensor fields: higher analogues of the electromagnetic field to which an electron couples.

My goal then-and still now, although I think we are finally getting to the bottom of it-was to understand what these higher analogues of the electromagnetic field actually are. Our first guess was what you just described: what mathematicians call a connection on a higher principal bundle. One chooses a higher analogue of the gauge group, such as the circle group U(1) in Maxwell theory. Here, “higher” does not refer to the dimension of the group, but to its conceptual stratification.

These higher principal connections certainly exist as mathematical objects, and I investigated them in my early articles. Eventually, however, it turned out that they are not generally the right description for the higher gauge fields that actually appear in string theory. There are applications in the heterotic sector, but the situation needs qualification.

I may be getting ahead of myself. Perhaps you wanted a simpler answer.

Interviewer: No, that is completely fine. I think the audience wants the technicalities.

Urs Schreiber: Then let me return to the idea of generalising ordinary gauge groups to 2-groups.

When we generalise a structure, we often begin by describing the ordinary structure carefully enough that the generalisation becomes apparent. The trouble is that an ordinary structure can have several equivalent descriptions whose higher generalisations are no longer equivalent. That is what happens here.

One way to describe a gauge field is through its parallel transport. This plays a central role in lattice gauge theory. In a background gauge field-whether the Maxwell field or the strong nuclear force field-a path traversed by a charged particle is associated with a group element. On a lattice, one assigns a group element to each edge, and this assignment describes the gauge field configuration.

More generally, without a lattice, we can consider suitably well-behaved continuous paths. A connection on a principal bundle, which is what physicists call a gauge field, assigns group elements to paths subject to compatibility laws. If two paths are composed, the group elements assigned to them are multiplied to give the element assigned to the composite path. This is parallel transport; around a closed path, it is holonomy.

Ordinary gauge fields are entirely characterised by their parallel transport. Physics textbooks usually define them using local gauge potential 1-forms, often denoted A, on charts of spacetime, glued together on overlaps. That is a good definition. But an equivalent description uses what physicists call Wilson lines and Wilson loops: knowing all of these observables amounts to knowing the gauge field. This is classical differential geometry.

The original idea was to extend this to a higher gauge field, such as the B-field in string theory, which is one step higher than the electromagnetic field. Instead of Wilson lines, we would use Wilson surfaces. To every little surface-for example, a piece of the worldsheet swept out by a string moving through spacetime-we would assign a group element, compatibly with gluing surfaces together.

In the abelian case, such as the B-field, where the element assigned to a surface lies in U(1), this works. It gives an equivalent definition of the higher gauge field.

Let me emphasise something for physicists who might be wondering what remains to be understood: “We learnt about the B-field in our first string theory course. It is a 2-form.” The point is to make the field globally well-defined on topologically nontrivial spacetime. We are talking about what one might call IR-complete, or IR-completed, higher gauge theories: theories whose field content is known globally, rather than just on a chart. Global topology introduces additional data and subtleties that a local differential form does not capture.

A major problem appears when we try to characterise higher gauge fields through higher parallel transport in the nonabelian case. For instance, in heterotic string theory, the B-field is combined with an ordinary nonabelian gauge field. Requiring well-defined surface transport then imposes a strong restriction on the admissible gauge field configurations: the so-called fake-flatness constraint.

Essentially, part of the higher gauge field must be flat. It is not the 3-form curvature that is required to vanish, but an associated 2-form curvature. This is an interesting constraint in many applications, but it is not the right constraint for the string-theoretic applications that originally motivated me and others.

That is why I eventually abandoned this particular description. More precisely, I abandoned the requirement that one have a principal higher connection of this kind: the idea that a higher gauge field must necessarily be completely characterised by its higher-dimensional parallel transport.

3. Where higher parallel transport works · 11:05

Interviewer: So, in this model, strings trace out surfaces. Are you saying the idea would not work for the B-field, or for fields such as electromagnetism and the strong force?

Urs Schreiber: No, let me clarify. It works for all ordinary gauge fields coupled to particles, including Maxwell theory and the strong force. It works for ordinary nonabelian gauge groups, and it also works for abelian higher gauge groups.

In particular, it works for the abelian B-field, the next higher version of the electromagnetic field, where an element of U(1) is assigned to each piece of surface. The difficulty arises when the gauge field is both higher and nonabelian.

If you ask what the RR-field in type IIA string theory is, or what the C-field in eleven-dimensional supergravity is, you can no longer simply say that it is described by a higher analogue of a gauge group whose elements are assigned to pieces of higher Wilson surfaces. You can still talk about higher analogues of gauge groups, but not with that particular characterisation by surface transport.

4. Why use toposes in physics? · 12:45

Interviewer: In relativistic physics there is no absolute spacetime, no great cosmic clock ticking away somewhere. People describe general relativity as background independent. Toposes, as I understand them, do not need an underlying topological space; mathematicians use sites instead. That seemed to me a reason to formulate physics in terms of toposes.

What initially led you to use toposes to answer physical questions?

Urs Schreiber: There is a claim in your introduction, and then a question. I would disagree with the claim: passing from topological spaces to toposes is not the analogue of passing from background dependence to background independence in general relativity. That is not what toposes do.

But I can explain what they do by answering your question. In this application, toposes are what mathematicians call context categories for generalised spaces.

Start by asking why quantum field theory is so much harder than quantum mechanics. A fair amount of the difficulty comes from dealing with spaces more general than ordinary manifolds. Let us look at one of the simplest ingredients.

In quantum field theory, a single configuration of the system is a field configuration. What is that? It is a map from spacetime to a target space. A complex scalar field is a map from spacetime to the complex numbers-a complex-valued function. A spinor field takes values in a spin representation. More generally, these are sections, but let us not complicate matters yet.

In quantum mechanics, the position of an electron is one point in a space. In field theory, the analogue of that point is an entire map. The configuration space is therefore a mapping space, and there are very many maps even from a fixed spacetime to the complex numbers.

In favourable cases these are infinite-dimensional manifolds, such as Fréchet manifolds. But for noncompact spacetime, as in physically realistic situations, the mapping spaces one needs need not carry the structure of an infinite-dimensional smooth manifold in the standard sense. Compactness makes the familiar constructions work; without it, the ordinary textbook account no longer supplies the space we need.

This means that already the configuration space on which we want to do physics-the space on which we would construct a symplectic form, a prequantum line bundle, and then sections for quantisation-is not available in the usual differential-geometric framework.

That may sound outrageous: “He is claiming that the space does not exist, yet everyone has been doing quantum field theory all along.” What usually happens is that one immediately leaves the full space behind and studies an infinitesimal neighbourhood of a classical solution. One does perturbation theory, and essentially everything becomes algebraic. This is where Feynman perturbation theory enters.

But if we want an IR-complete theory, globally defined, we need the full space. To go beyond the usual Feynman-Dyson-Tomonaga perturbative account, we need a home for the configuration spaces that occur in field theory.

That is where toposes enter. Suitable toposes are categories of generalised spaces in which these mapping spaces naturally exist as differential-geometric objects. We can genuinely ask whether they carry symplectic forms, whether we can construct prequantum line bundles, and whether we can quantise them.

This is what I try to explain in my encyclopaedia article, “Higher Topos Theory in Physics”: toposes provide the missing generalised spaces needed in physics.

And mapping spaces are only the beginning. Next come fermions, even before supersymmetry. Classical fermions are not ordinary functions on spacetime; they are odd, or Grassmann-odd, functions. Again, this is outside the standard elementary differential-geometric account, and topos theory supplies a natural context for their configuration spaces.

Then come gauge fields, higher gauge fields, higher gauge-of-gauge transformations, and perhaps orbifolds. Beyond perturbation theory, we need spaces accommodating the infinite-dimensional nature of fields, the odd grading of fermions, and the higher structure of gauge fields. These are difficult to handle together with traditional tools, but they all live naturally in higher topos theory.

5. The successes and limitations of perturbation theory · 20:25

Interviewer: Perturbation theory in QED famously agrees with experiment to something like eleven decimal places. But, as you have said, nature is nonperturbative: it works beyond an infinitesimal neighbourhood. Is there a deep reason why quantum field theory can give such precise results, yet the series blows up if you sum all the terms?

Urs Schreiber: The first point is that this success is not general. The often-repeated claim that quantum field theory is the most accurate theory humanity has ever devised refers specifically to quantum electrodynamics. That was the original, enormous success of quantum field theory, and we are still living in its echo.

It is true that QED works extraordinarily well. It is harder to pin down precisely why, but the intuitive reason is that QED is weakly coupled. The fine-structure constant is approximately 1/137, a small number, so one expects many successive loop orders to continue improving the approximation.

The trouble is that perturbation theory does not work nearly so well throughout the Standard Model. In QCD, for example, the theory is not weakly coupled at ordinary temperatures. Weak coupling emerges only at extremely high temperatures. To understand bound states such as hadrons, we need strongly coupled QCD, precisely where perturbation theory does not provide that extraordinary accuracy.

This matters when experiments at the LHC report an anomaly-a deviation from a Standard Model prediction. There can be uncertainty about whether the experiment really disagrees with the theory or whether the prediction suffers from hadronic uncertainties.

Even processes that do not nominally involve quarks or hadrons receive contributions from them through loops. If perturbation theory cannot compute those contributions, they must be supplied through other means: comparison with other experiments, computer simulations, or models that have proved useful. We do not obtain all of those contributions from the perturbative first-principles calculation itself.

Even in QED, the agreement with experiment is somewhat mysterious, because the mathematics of the perturbative procedure does not by itself justify it. After renormalising, cancelling anomalies, carrying out BV quantisation, and fixing the ghost sector, one finally computes a scattering amplitude. One expects a number to take to an experiment and compare with a measurement.

Instead, in the physically interesting theories, one obtains an asymptotic power series. It looks like a Taylor series, but it is not a convergent expression for a number. The series has zero radius of convergence.

This is a serious conceptual difficulty. In QED, one nevertheless adds the first few terms and finds that the result approaches the measured value with remarkable accuracy. Rules of thumb suggest that, because the coupling is about 1/137, one might continue for a substantial number of orders before the divergence becomes apparent. But this is an intuitive account. The mathematical output is a nonconvergent series, and one truncates it and treats that truncation as a prediction.

To illustrate the logical problem, imagine that I claim to have a theory of everything consisting of the harmonic series, 1 + 1/2 + 1/3 + ..., whose positive terms decrease while its partial sums grow without bound. Give me a number-your telephone number, or a stock price-and I keep summing until I get close to it, then stop. Because the partial sums grow without bound, they eventually pass any chosen positive number. Because the increments decrease, larger target numbers can be approached with increasing precision.

That is deliberately silly. But it draws attention to the question of what justifies stopping a divergent series at a particular point and calling the result a prediction.

Of course, people know this problem. They introduce resummation procedures, such as Borel resummation. One can do that, but it is essential to recognise that this supplies further input. We are supplementing the original theory. In effect, we are acknowledging that the initial account was incomplete and needs additional structure.

Interviewer: So that is part of what makes quantum field theory hard: despite the extraordinary accuracy in some cases, more input is needed to give the perturbative expressions a meaning.

6. Generalised spaces through their probes · 28:24

Interviewer: Suppose someone is reading Lawvere or Emily Riehl and encounters these definitions. They might ask: “How do I probe a topos?” How does one actually extract information from it?

Urs Schreiber: Let me first clarify the question. In this physical application, it is usually the objects in the topos that we want to probe, rather than the topos itself.

I claimed that topos theory gives us access to generalised phase spaces that ordinary differential geometry does not provide. How does it do that? The language of probes offers a good explanation. Here I am speaking about Grothendieck toposes, as distinct from elementary toposes.

Traditionally, we define a space by beginning with a set of points and then adding structure. We specify open neighbourhoods to obtain a topological space, and add smooth structure to obtain a manifold. The points come first.

Topos theory allows a more operational approach, closer in spirit to how we physically encounter space. I do not know the space around me because I carry a list of its points in my pocket. I send light through it and see what returns, or reach out with my hand to see what is there. I probe it.

Mathematically, we begin by agreeing on a collection of simple test spaces-for example, small n-dimensional balls, for all n. I then describe a space by telling you, for every test space, the set of ways that test space can map into it.

This sounds circular: the space has not yet been defined, and I am already telling you about maps into it. Think of it as a bootstrap procedure. I supply its operational behaviour under all probes.

These assignments must be consistent. Suppose a smaller test ball maps into a larger one. Every probe from the larger ball into the space then gives a probe from the smaller ball, by composition. In technical language, probes pull back along maps of test spaces, and this pullback respects composition and identities.

Such data-a set for each test object, together with compatible pullback maps-form what mathematicians call a presheaf of sets.

We add one further ingredient: locality. Suppose a large test region is covered by smaller regions. A probe from the large region restricts to probes on each smaller region, and those restrictions agree on overlaps. We require the converse: compatible local probes must glue uniquely to a probe from the whole region.

In other words, if a large disc is covered by little discs, probing with all the little discs should suffice, provided the probes agree where the discs overlap. This is the sheaf condition. A presheaf satisfying it is a sheaf, and a Grothendieck topos is a category of such sheaves on a site.

The site is our collection of test spaces, with the maps and covering information that make this description possible. It specifies the probes with which our generalised spaces are modelled.

You might still object that I have not said what the space itself is. I have only described a consistent system of probes. The crucial move is to regard that system-the sheaf-as the space itself.

Now we can do geometry. Given two such systems of probes, a map between them should carry probes of the first to probes of the second, compatibly with all the structure. This produces a category of generalised spaces.

There is an especially important consistency check. Our original test spaces now have two incarnations: the ordinary spaces we began with, and the generalised spaces given by their own systems of probes. The Yoneda lemma tells us that these descriptions faithfully recover the original test spaces and their maps. They embed faithfully into the new world of generalised spaces.

7. Relational quantum mechanics and structuralism · 37:16

Interviewer: You described probing an object by bouncing a photon off it. Carlo Rovelli and others have proposed relational quantum mechanics: rather than simply asking what is in the Andromeda galaxy independently of any interaction, one considers photons or other interactions that connect it to us. Is there a connection?

Urs Schreiber: At a philosophical level, there is a similarity. Mathematically, however, these are not the same thing. A topos by itself does not know about quantum mechanics, let alone relational quantum mechanics.

There is a broader connection to what philosophers of mathematics call structuralism. This is sometimes contrasted with the material perspective of classical set theory, in which a set consists of definite elements and those elements determine its identity.

In category theory, objects acquire their meaning through the maps between them. That is precisely what we have just used: we described generalised spaces without opening them up and inspecting an underlying collection of points. We looked at how other objects map into them.

Independently of the physical application, this can be developed into a foundation of mathematics. It is called structuralism because the nature of mathematical objects is expressed through their structural relations. In that sense, the philosophy resembles the one motivating relational observables in quantum gravity.

8. Lawvere and synthetic differential geometry · 39:33

Interviewer: Who first used topos theory in physics? Was it Lawvere, Chris Isham, Andreas Döring? What earlier work does this build on?

Urs Schreiber: There is an interesting story here. I think a case can be made for Lawvere, who also essentially invented the relevant form of topos theory. His physical motivation is not widely known, and it is perhaps more visible in his unpublished notes and interviews than in his published articles.

Apparently, as a student, he found himself in a physics department and was asked to work on a problem in continuum mechanics involving vector fields and their flows. He stopped to ask: “What is a vector field?” One can imagine the distress of an adviser who had hoped he would solve the assigned problem! Everyone else regarded the meaning of a vector field as settled.

From his interviews, it appears that this question helped motivate his development of synthetic differential geometry, a topos-theoretic approach to differential geometry.

The idea can be made quite concrete. A standard mathematical definition of a tangent vector uses an equivalence class of smooth paths through a point: paths are equivalent when they have the same derivative there. This is precise and rigorous, but it is not how a physicist usually imagines a vector.

A physicist imagines a very short directed line-an infinitesimal displacement. A force points a little way in a certain direction. We begin at a point and move an epsilon in that direction. Intuitively, a vector should literally be an infinitesimally short path.

That intuition goes a long way in physics, but giving it a rigorous mathematical form requires something more. Related ideas had appeared in algebraic geometry through nonreduced schemes, and Lawvere transported them into this setting. He developed elementary topos theory and axioms for synthetic differential geometry.

Remember that toposes are context categories for generalised spaces. Some contain infinitesimal spaces: an interval that is not simply short, but infinitesimal, without collapsing to a point. A tangent vector in a manifold can then literally be a map from this infinitesimal interval into the manifold.

That is the sense in which the geometry is synthetic. The infinitesimal object is available directly. We do not first construct the vector from equivalence classes of paths and derivatives; it is an actual infinitesimal path in the space.

Perhaps it is a stretch to say Lawvere was the first to apply topos theory to physics. He was apparently motivated by physics and went on to develop elementary topos theory, internal languages, and synthetic differential geometry. I am not sure whether he returned full circle to answer the original physical question.

Interviewer: I think he wrote a book on continuum mechanics and category theory, if I remember correctly.

9. Singularities and quantum measurement · 44:24

Interviewer: On Curt Jaimungal's *Theories of Everything* channel, you mentioned that physicists sometimes give things rather silly names, such as "Big Bang" or "black hole." Do toposes help with singularities, infinite curvature, or other historically difficult phenomena, perhaps even wavefunction collapse?

Urs Schreiber: I would not say that these are problems automatically solved by toposes. General relativity lives perfectly happily in a topos of synthetic differential geometry, but that does not itself resolve the issues you mentioned.

One relevant example is an orbifold singularity. Such singularities arise when a space looks like a quotient by a group action that is not free everywhere. Imagine taking a flat piece of paper and forming a cone. At its tip, the result is no longer an ordinary smooth manifold. Quotient constructions of this kind lead to orbifolds.

Higher topos theory provides a very natural home for these singularities. My books *Equivariant Principal Infinity-Bundles* and *Geometric Orbifold Cohomology* are about precisely this kind of geometry.

Wavefunction collapse is a separate question; it is not, at least a priori, a question about black hole singularities. But it does bring us back to Isham and Döring, whom you mentioned earlier.

The idea is to make sense of the logical nature of quantum physics, including the apparent collapse of the wavefunction on measurement. People find this mysterious and sometimes speculate about consciousness. A mathematical approach would instead ask whether we need a different logic.

Could we find a logic in which the ordinary description of classical measurement, when interpreted there, becomes the quantum measurement process? Then we could say that quantum measurement is structurally like classical measurement, but seen internally to a different topos.

This is another major aspect of topos theory: it provides nonstandard logics. Peter Johnstone's famous book is called *Sketches of an Elephant* because topos theory has several aspects, like the elephant in the story of the blind men. From one end it seems like a snake; at a leg it seems like a tree trunk. We have talked about geometry, but there is also this logical aspect.

Within a topos, objects can be interpreted not only as spaces but also through their logical meaning. Isham, and subsequently Döring, pursued the idea of constructing a topos whose internal logic would express quantum mechanics. Once immersed in that mathematical universe, ordinary-looking logical statements would become statements of quantum logic. Perhaps quantum phenomena seem strange partly because we view them through an inappropriate logical framework.

A group in the Netherlands developed a related construction, which I think worked better for that purpose, known as the Bohr topos approach. It seeks to formalise a dictum associated with Niels Bohr: whatever happens in the quantum realm must ultimately be expressible in classical terms.

This brings us back to probes. We take classical contexts as our test objects-technically, commutative C*-algebras-and describe the quantum object through its classical probes. That is one strategy for making quantum phenomena more intelligible at the logical level.

10. Why the internal logic is not generally classical · 51:12

Interviewer: Some listeners may have heard of intuitionistic logic and the failure of the law of excluded middle. People sometimes say that machines think in Boolean logic while humans do not, but people also have difficulty imagining a cat that is neither dead nor alive. We often think in Boolean terms ourselves.

What kinds of logic live inside toposes? What are their properties? Perhaps people have heard of Heyting algebras.

Urs Schreiber: The internal logic of a topos is famously intuitionistic, or constructive. The law of excluded middle does not hold in general, and neither does the axiom of choice.

The axiom of choice offers a useful example because there is a geometric way to see why physicists might want it to fail internally.

In its usual formulation, we have an index set and a nonempty set over each index. For example, for every natural number n, we have a set S_n. Choice says that we can select an element from each S_n. With natural numbers as indices this is countable choice; the general axiom allows other index sets.

Now rephrase this geometrically. Regard the index set as a base space X. Put the corresponding set over each point as a fibre. The disjoint union of those fibres is a total space P, with a projection P → X that sends every element of S_n to n.

A choice of one element in each fibre is exactly a section: a map X → P whose composite with the projection is the identity on X. In this language, the axiom of choice says that every surjective map of sets has a section.

Now suppose we are in a context where our objects are not bare collections of points, but spaces with topology or smooth structure. The projection is a continuous or smooth map, and a section must also be continuous or smooth. Nearby points downstairs have to lift compatibly to nearby points upstairs.

At that point, the failure of choice becomes geometrically natural. A central feature of bundle theory is that bundles can have global obstructions to sections. Even if ordinary set-theoretic choice supplies a selection of points, that selection need not be continuous or smooth.

If every relevant bundle in smooth geometry had a smooth global section, the nontrivial topological sectors that interest us in physics would disappear. Nontrivial charge sectors and the kinds of global phenomena encoded by these bundles rely on obstructions to such sections.

So the nonclassical logic is not merely an abstract oddity. A topos is a category of spaces with particular geometric structure. Its objects are not plain sets, and the ordinary logic of plain sets need not apply internally. Does that help?

Interviewer: Yes.

11. Topological phases of matter · 57:45

Interviewer: I would like to turn to my favourite subject: condensed matter systems. You mentioned recent work in condensed matter physics. What do toposes have to say about many-body systems?

Urs Schreiber: This takes us back to perturbation theory. The application is to the topological sector of many-body systems, especially topological phases of matter.

These are solid-state systems whose important properties may be visible only globally. Locally, a system may look like another system, while globally there is a twist that distinguishes them.

At a broad level, this resembles the problem of bound states in QCD. Perturbation theory examines a neighbourhood of a point in configuration space. It sees the nearby behaviour but can miss global structure. Hadronic bound states are not visible in ordinary perturbation theory; similarly, topological phases of matter involve structure that is nonperturbative in this global sense.

This is why methods we developed over the years, originally motivated by high-energy physics and string theory, turned out to apply to topological phases of matter. One particularly successful application concerns fractional quantum Hall systems, where these global effects are experimentally observed as topological order.

12. Fractional charge and the quantum Hall effect · 1:00:25

Interviewer: In one of your videos, you gave a very clear visual explanation of filling factors and fractional charges in terms of flux quanta and the number of electrons. Could you explain how fractional charges arise, and how toposes enter the description of the fractional quantum Hall effect?

Urs Schreiber: Let us begin with the physical effect itself. At this stage, no sophisticated mathematics is required; it is an experimental observation.

We confine electrons essentially to a two-dimensional surface, for example at an interface between two semiconductors, creating an electron gas that can explore only that surface. We then apply a strong magnetic field perpendicular to it.

Magnetic flux quantisation, which is itself related to global topological structure, gives us flux quanta. For the basic fractional quantum Hall states we are discussing, we arrange for a particular integer number k of flux quanta per electron.

Because there are many electrons, this requires a large amount of magnetic flux-fields of several tesla. The ratio also has to be finely controlled. In fact, the required behaviour would be difficult to achieve simply by tuning an experimental knob were it not for localisation effects: disorder in the system helps to stabilise the situation. But let us leave that mechanism aside for now.

Interviewer: Could we give that a concrete numerical example?

13. Two electrons, seven flux quanta, and a fractional quasihole · 1:03:07

Urs Schreiber: The simplest, historically important example has three flux quanta per electron: the filling fraction $\nu = 1/3$.

Interviewer: So if I have two electrons and seven flux quanta, six of the flux quanta give three per electron, and the extra one produces a quasiparticle with one-third of an electron’s charge?

Urs Schreiber: That is the idea. The state with exactly k flux quanta per electron is the fractional quantum Hall state, with its characteristic experimental signatures.

Something happens in this state that is not easy to describe theoretically. The electrons interact strongly through the Coulomb force. They are no longer weakly coupled, and their collective behaviour produces new phenomena. We can observe these in experiments even when a complete nonperturbative theoretical account is difficult.

Quasiparticles appear when we move slightly away from the exact state. Starting with k flux quanta per electron, we can add one flux quantum or one electron. Suppose we add a flux quantum. Against the background in which everything was paired at the ratio k to one, that extra flux behaves as though 1/k of an electron were missing.

In your example, k is three: two electrons correspond to six flux quanta, and the seventh produces a quasihole behaving as a deficit of one-third of an electron.

Interviewer: Your pictures made this much easier for me to understand. I recommend that people watch that video. I had struggled with filling factors before seeing it.

Urs Schreiber: The talk is called “The $1 Million Puzzle: Hadrons and Quantum Computers - Why Modern Physics Needs a New Global Mathematics”.

Interviewer: You have your own YouTube channel, which I will link below. Did you record that presentation twice?

Urs Schreiber: At least twice. There was no substantial difference in content. One version was an evening lecture in German, with the same graphics; another was recorded outdoors.

Interviewer: I heard something in the background and thought you were on holiday, perhaps in Southeast Asia.

Urs Schreiber: No, I was at work in the Emirates. I recorded it while walking around the palm garden at a university in Sharjah. I like to walk, or at least stand, when giving these talks, so I take my laptop along and record as I walk.

Interviewer: I thought you were using your phone, but of course you needed your laptop to show the slides. The sound was quite clear.

14. What IR completion means · 1:07:14

Interviewer: Graduate students may be familiar with UV completion: extending a low-energy effective theory to higher energies. You speak instead about IR completion. Why are you looking for that?

Urs Schreiber: Again, this is closely connected with the nonperturbative story. UV completion gets a great deal of attention. People discuss the need for a UV-complete theory of quantum gravity, or M-theory as a UV completion of superstring theory. These are worthwhile questions, but completing a physical theory involves more than its behaviour at very short distances.

“UV” and “IR” are physics jargon. Ultraviolet light has short wavelengths, so UV refers to small scales. Infrared light has long wavelengths, so IR refers to large scales and long distances.

IR completion means supplying the missing definitions and specifications that make a field theory well-defined at large scales-ultimately, on the largest scales of the entire system. Just as UV completion asks for the theory at the smallest conceivable scales, IR completion asks for its global definition. That brings us back to topological twists.

Consider electromagnetism historically. Faraday and Maxwell developed the electromagnetic field, which we now describe using a differential 2-form. We can locally encode it through a gauge potential 1-form A.

Dirac’s insight in 1931 showed why this local account is not the whole story. On topologically nontrivial spacetime, a gauge potential need not exist as a single globally defined 1-form. It may exist on each contractible chart, but not over the whole space.

The modern description is that the electromagnetic field is a connection on a principal bundle. Local gauge potentials are glued by transition functions. Those functions encode the global bundle structure and make the field well-defined throughout spacetime, potentially across the cosmos.

In this sense, the global structure associated with Dirac charge quantisation is an IR completion-one of several possible choices of IR completion-of the local Maxwell theory.

The original insight is nearly a century old. Yet the analogous question for other gauge fields, especially higher gauge fields, has received much less attention. That is one of the questions I have been working on.

Only after specifying such an IR completion do we know the global topological charges of the theory, including anyonic charges or, in higher dimensions, brane charges.

Dirac’s magnetic monopoles provide the basic example. A field sourced by a magnetic charge requires globally nontrivial topological structure. Conversely, once we know the field’s global structure, we know which topological charges it can carry.

So before discussing renormalisation and interactions at tiny distances, there is another question: what is the actual global definition of the fields, even classically? That definition is what lets us discuss their topological charges.

15. Global structure and topological protection · 1:13:00

Interviewer: In quantum computing, would an IR completion lead to long-range topological protection?

Urs Schreiber: That is exactly the idea. Topological effects in solid-state physics are global effects. If we write an effective field theory for a material and want to understand its topological phases, we need the IR-complete theory.

“Large” here means large relative to the system: the complete sample may still fit in your hand, or be a tiny specimen in a laboratory. The point is that there may be a topological twist across its full extent.

16. Anyons and quantum materials · 1:14:04

Interviewer: You also use the phrase “geometric engineering.” Once we have an IR-complete theory, would engineers manufacture something with the nontrivial topology specified by the theory?

Urs Schreiber: Yes, although the interesting situation is that experimentalists already know how to produce these effects. It is often the theorists who are lagging behind in understanding them.

The fractional quantum Hall effect, for example, has topological order and anyonic phenomena. It has existed in the laboratory for decades. The challenge for theorists is to understand those topological structures through an appropriate effective field theory.

Interviewer: Could you explain anyons, their statistics, and how they appear in fractional quantum Hall materials?

Urs Schreiber: We discussed quasiholes that appear when an additional flux quantum is introduced into a state with filling fraction 1/k. Each behaves like a missing fraction 1/k of an electron.

Now take two such quasiholes and arrange for them to move around each other. The system is confined to a two-dimensional surface. The many electrons form one strongly interacting quantum system with a collective quantum state.

When the quasiholes go around each other and return to their original positions, everything appears to be back where it started, but the quantum state has acquired a complex phase. That phase can be measured. These effects have been observed experimentally in recent years, since 2020.

Interviewer: So their worldlines form a braid?

Urs Schreiber: Exactly. Viewed as a process in spacetime, moving the quasiholes around one another traces out a braid.

Interviewer: And we do not know how to do this in three dimensions?

Urs Schreiber: We can describe three-dimensional analogues theoretically, but it is unclear how to realise them experimentally with current methods. One could imagine a three-dimensional material containing little loop defects instead of point particles. The loops could turn around one another or thread through one another, giving a richer range of motions.

On paper, one can conceive of topologically ordered systems in which the collective quantum state acquires phases as these loops move. Experimentally, however, one first has to create and control suitable loop defects, which is difficult.

17. Classifying spaces and the choice of a global theory · 1:17:53

Interviewer: Could you explain classifying spaces, perhaps starting with electromagnetism?

Urs Schreiber: I have been emphasising a gap in the traditional treatment of nonperturbative, IR-complete phenomena. We had individual examples, such as electromagnetism, but lacked a general account.

If we want to specify an IR completion of a higher gauge field, the mathematically powerful way to do it is through a classifying space.

Earlier we described a field as a smooth map from spacetime to a target such as the complex numbers or a spin representation. A topological charge can also be encoded by a map. Think of it as an additional field component, visible at the global, IR level: a continuous map from spacetime into another target, regarded now as a topological space.

This target is not physical space. It may be quite abstract, or relatively easy to visualise. What matters is its topology. Topological charges are encoded by homotopy classes-deformation classes-of maps into it.

There are two claims here. First, the charges of an IR-completed theory can be encoded this way. Second, for given local higher gauge field laws, there is a way to determine which classifying spaces are admissible as flux quantisation laws.

Interviewer: But the classifying space is not unique? There can be several choices?

Urs Schreiber: Exactly. That is essential. The classifying space specifies the completion, or flux quantisation, of the higher gauge theory. A given local theory generally admits infinitely many possible choices. IR completion requires additional physical input; it is not automatic.

This is an underappreciated point. A Lagrangian density usually specifies a theory only chart by chart, since it contains gauge potentials that need not be globally defined. We still have to say which global theory we mean.

Even in electromagnetism, people tend to treat Dirac charge quantisation as the final, uniquely prescribed IR completion. But there is room for alternatives, and eventually one should examine which flux quantisation laws are selected by experiment.

Magnetic monopoles have famously not been found. Predictions about them therefore do not, by themselves, settle the choice. On the other hand, related flux quantisation is visible in Abrikosov vortices in type-II superconductors, where magnetic flux threads the material in quantised vortices. These have been observed since the 1960s.

That tells us the familiar Dirac completion is at least on the right track: its integer flux quantisation is experimentally visible. There could nevertheless be modifications that change other details not yet distinguished experimentally.

So even the statement that the electromagnetic field is precisely a connection on a principal circle bundle contains a choice that deserves scrutiny. There may not be dramatic surprises in this simplest example, but the freedom becomes larger and more interesting for higher gauge theories.

18. From electromagnetism to a model of the quantum Hall state · 1:24:10

Interviewer: Is the electromagnetic field the field that quantum computing experimentalists would work with here?

Urs Schreiber: More precisely, it is what becomes of magnetic flux in the fractional quantum Hall state.

At the special filling fraction, the electrons and magnetic flux form a strongly coupled collective state. Extra flux quanta behave as quasiholes or anyons. We are no longer examining magnetic flux against the ordinary vacuum, but against this very particular electron liquid.

We want an IR-complete effective field theory of that surplus flux relative to the strongly coupled background. Its effective description need not be the same as that of a magnetic field in empty space. This suggests looking for a modified classifying space.

Interviewer: So we have the electron gas in that special configuration, and we seek a classifying space appropriate to it?

Urs Schreiber: Yes. The classifying space for ordinary electromagnetism with Dirac charge quantisation is CP^∞, infinite-dimensional complex projective space.

It can be built in stages. We begin with a 2-sphere, then attach a 4-dimensional cell, then a 6-dimensional cell, and so on. Informally, one can picture taking a higher-dimensional sphere, removing a piece to leave a cell, and attaching its boundary to what has already been constructed.

Interviewer: Is a familiar fibration involved?

Urs Schreiber: Yes. At the first nontrivial step, the attaching map from the boundary 3-sphere to the 2-sphere is the Hopf fibration. Continuing through the higher even-dimensional cells produces CP^∞.

This is an abstract mathematical space, not somewhere one can walk or send a rocket. It classifies ordinary electromagnetic charge configurations.

For the quantum Hall system, we look for a related but different classifying space. We know the surplus flux does not behave like free flux in a vacuum: braiding it produces a special response. A modification of the classifying space should reflect that difference.

Interviewer: Is it easier to find the classifying space in a vacuum than in an electron background?

Urs Schreiber: I would not necessarily say easier. The vacuum story is simply the one that has been known for decades. The fractional quantum Hall system is a new effective system, unknown to Maxwell, so we must propose an effective description for it.

The construction of CP^∞ suggests a simple guess: stop at a finite stage instead of continuing indefinitely. The simplest possibility is the very first stage, CP^1, which is the 2-sphere S^2.

We can therefore ask what happens if we flux-quantise Maxwell theory-or, more precisely, Maxwell-Chern-Simons theory-using S^2 as the classifying space of charges.

Once that choice is made, the IR charge sector is specified. For this topological sector, we do not need a Lagrangian density: the configurations are classified by homotopy classes of maps from the spacetime inhabited by the electron liquid into S^2. We can then analyse the homotopy type of the mapping space mathematically.

The remarkable result is that this choice-flux quantisation in degree-two cohomotopy-recovers, in considerable detail, known and expected properties of fractional quantum Hall systems. That is evidence that it is a useful and possibly correct choice.

19. What the model predicts for quantum hardware · 1:30:53

Interviewer: Suppose S^2 is a good choice and I work out the corresponding classification. What does that tell me about manipulating a fractional quantum Hall liquid or material? What extra information does it provide?

Urs Schreiber: It works like any physical theory. We propose a formalism, check that it agrees with known observations and established expectations, and then ask what else it predicts. If the new predictions fail, we revise the model. If they succeed, we continue.

Here, the model reproduces a substantial amount of what is already known about fractional quantum Hall systems, which encourages us to investigate further.

We recently published an article making a new prediction of this kind: "FQH Liquids with Flux-Expulsion Islands allow Nonabelian Anyons" (DOI).

Consider an electron liquid containing small islands from which the surplus flux is expelled. One possible mechanism would be superconducting islands within the semiconductor, expelling flux through the Meissner effect. Another mechanism might work too, but superconductivity is the one I have in mind.

We can take the effective theory with its cohomotopy flux quantisation and calculate how these islands change the topological properties. Something striking happens: the theory predicts another form of anyonic behaviour, now associated with the islands rather than with freely moving quasiparticles.

With three or more such islands, the allowed transformations of the quantum state can become nonabelian. Instead of multiplication by a complex phase, the state can transform by higher-rank unitary matrices. That is the sort of structure required for useful topological quantum computation.

Abelian phases are interesting and have been observed, but on their own they do not supply universal topological quantum gates. Nonabelian transformations offer a route to a sufficiently rich set of operations.

Thus the IR-completed effective theory of surplus flux predicts a potential new route to nonabelian anyons in fractional quantum Hall systems.

Two points make this especially interesting. First, there are many proposals for anyonic systems, but fractional quantum Hall systems are the setting in which the relevant anyonic phenomena have been repeatedly observed by independent experimental groups. Other proposed platforms are either still theoretical or lack that level of independent reproduction. This makes fractional quantum Hall matter a particularly concrete candidate for topologically ordered quantum hardware.

Second, people have already investigated how to obtain nonabelian anyons in these systems. The famous candidate is the filling fraction 5/2, for which there are theoretical arguments and encouraging experimental indications of nonabelian behaviour. But experimentally that state is extremely fragile and requires exceptionally low temperatures. Its fragility has long raised doubts about its practical use for stable quantum computing.

The attraction of the proposal with superconducting islands is that it could provide a more stable form of nonabelian behaviour. That remains a prediction and a potential route to hardware.

20. Why nonabelian operations matter · 1:36:45

Interviewer: Some listeners will know the Aharonov-Bohm effect, in which an electron acquires an observable phase. That phase is a complex number. For abelian anyons, we similarly have a complex phase, whereas the nonabelian analogue is a matrix?

Urs Schreiber: Yes, a unitary matrix.

Interviewer: And that allows a universal set of quantum computing gates?

Urs Schreiber: That is the goal. In quantum computing, we want a collection of elementary processes that can be implemented reliably on hardware and combined to approximate any desired quantum computation.

Classical computers work this way. They have simple logical gates acting on a few bits. Everything complicated that we do, including this video conversation, is built from enormous combinations of those elementary operations. The elementary set must be universal: sufficiently expressive to build all the computations we need.

For quantum computation, the corresponding requirement is that any desired unitary operator can be approximated, within a specified error, by composing and tensoring elementary unitary operators. Those elementary operators are the gates.

If every gate merely multiplies by a complex number, then no amount of composition or tensoring changes that essential limitation. Tensor products of one-dimensional Hilbert spaces remain one-dimensional. Such operations alone cannot be universal. Higher-rank unitary matrices are needed, although determining which collections are universal is a further question.

The long-term aim is to go beyond NISQ devices-the current noisy intermediate-scale quantum devices-and find topologically protected processes that provide a universal set of quantum operations.

21. Quantum programming languages and linear type theory · 1:39:45

Interviewer: You also work on quantum programming languages. I have read that classical programming languages are related to cartesian closed categories. What changes for quantum languages? Would we need homotopy type theory, or linear homotopy type theory? I am not familiar with type theory myself.

Urs Schreiber: Let me start with the distinction concerning products. Quantum systems require a noncartesian tensor product, and the corresponding categorical setting is a suitable closed monoidal category.

This is closely connected with entanglement. At a simple classical level, a system has a set of states. Combining two systems gives the cartesian product of their state sets. A state of the composite is just a pair: one state of the first system and one of the second.

For quantum systems, the state spaces are Hilbert spaces. Let us temporarily ignore the extra Hilbert-space structure and think just of vector spaces. The composite system is described by their tensor product, not their cartesian product.

The tensor product is another way of combining two objects into an object of the same kind, but it does not satisfy the defining properties of a cartesian product. This noncartesian nature is the mathematical reflection of entanglement.

An element of a tensor product need not be a single pair of component states. It may be a superposition of such product states. Much of the distinctive behaviour of quantum mechanics, and much of the potential power of quantum computing, comes from this fact.

Recall our discussion of a topos having an internal logic. An ordinary topos is cartesian closed, so its cartesian product by itself does not model quantum composition. We need additional structure, a tensor product describing the combination of quantum systems.

The precise language for discussing these internal logical structures is type theory. Think of an object as a data type, and a map between objects as a program that accepts data of one type and returns data of another. This gives a programming-language perspective on the category.

To describe quantum phenomena, we need a suitable linear type theory. In the physical interpretation, coherent quantum processes are described by linear maps between state spaces, rather than arbitrary functions. The type theory must reflect that structure.

22. Bringing homotopy theory into the picture · 1:45:10

Interviewer: Where does homotopy enter linear homotopy type theory?

Urs Schreiber: This takes us into deeper waters, but there is a fascinating story.

You may have noticed a tension in what I said. Toposes have internal logics, and we would like those logics to account for quantum phenomena. Yet their native product is cartesian. In the sense relevant here, toposes are classical objects. Indeed, I first motivated them through classical configuration and phase spaces: maps from spacetime into field targets, equipped with symplectic structure.

An ordinary topos is excellent for this classical part of physics, but it does not automatically contain the linear structure and tensor product needed for quantum processes.

Categories of vector spaces belong to the world of abelian categories. At the level of ordinary category theory, toposes and abelian categories are like oil and water: their overlap is only trivial. We cannot simply merge the two structures in the most direct way.

Something remarkable happens on passing to higher category theory, specifically (∞,1)-category theory. Ordinary toposes become infinity-toposes, while the higher counterpart of the linear setting is provided by stable infinity-categories.

These are still different kinds of structure, but they begin to fit together in a new way. In suitable presentable cases, both can be approached through presheaf constructions.

Now imagine a generalised space, with its higher gauge symmetries, carrying an additional linear structure over it. As an intuitive picture, think of a space equipped with a bundle of vector spaces or Hilbert spaces-something reminiscent of a prequantum bundle. Technically, the relevant higher objects are homotopy types equipped with parametrised spectra.

A remarkable theorem, first noticed by Georg Biedermann in private communication with André Joyal in 2007 and first published by Joyal in §35 of *Notes on Logoi* (2008), says that the category of these parametrised spectra is itself an infinity-topos. Further background is available in the nLab entry on Joyal loci. At first this seems unlikely: the objects contain linear information, with tensor products and quantum-like structure. But that information is not considered in isolation. It is parametrised over a classical space, and the combined objects again form an infinity-topos.

These are tangent infinity-toposes. They retain the advantages of higher topos theory and its internal logic while containing the linear structures needed for quantum phenomena.

This is suggestive in relation to Bohr's idea that quantum physics must be accessible through classical contexts. Here we have a mathematical setting in which the quantum, linear information is carried over classical parameters.

23. A classical context containing quantum information · 1:51:04

Interviewer: So the Grothendieck toposes you mentioned describe classical physics. To include quantum effects, you thicken the objects with these quantum bundles and pass to an infinity-topos, perhaps a cohesive infinity-topos?

Urs Schreiber: Cohesion is not needed for the statement at this point. The point is that the higher versions of toposes and linear categories can be combined in this particular way.

The result is again an infinity-topos: a classical kind of context containing quantum information internally. It is another way to think about Bohr's dictum. We have a classical universe with a means of accessing quantum information inside it.

If we isolate the linear quantum information, we are in a stable infinity-category, analogous to an abelian category, rather than a topos. But when it is carried over classical parameters, the combined setting can provide the richer logical framework of an infinity-topos. That is probably the best I can do in plain language; ultimately, it is an abstract mathematical statement.

24. Homotopy, gauge equivalence, and quantum programming · 1:52:47

Interviewer: Where exactly does homotopy enter the theorems you mentioned, and the construction of a quantum programming language for these infinity-toposes?

Urs Schreiber: It is surprising, because homotopy does not initially seem to have anything to do with quantum theory.

One useful way to think of homotopy theory is as the mathematician’s language for the gauge principle. In physics, we often should not ask whether two field configurations are literally equal. They may look different while being gauge-equivalent. More precisely, we should ask for a gauge transformation relating them. That is what the homotopy theorist describes as a homotopy.

In higher gauge theory, the same question recurs. Two gauge transformations between the same configurations need not be literally equal; there may be a gauge-of-gauge transformation between them. Then there may be transformations between those transformations, and so on.

Taken to its logical conclusion, the gauge principle says that literal equality is often the wrong question. We need equivalence, together with the transformations witnessing it, at every level.

Higher category theory provides the modern mathematical language for this structure. An ordinary category describes objects through their maps. An infinity-category adds higher transformations between those maps, transformations between transformations, and so forth. This is the natural setting for gauge fields and their higher symmetries.

Correspondingly, an infinity-topos is a context for generalised spaces with higher gauge structure built in. Since ordinary toposes have internal type theories, we can ask for the internal programming language of these higher toposes. That is homotopy type theory.

Now recall the special infinity-toposes we discussed, the tangent infinity-toposes, which also contain linear structure. We should expect a version of homotopy type theory that reflects that additional structure: linear homotopy type theory.

I raised this question around 2014 and pointed out structural indications that such a theory should exist and behave well. Several years later, perhaps four or five years ago, someone wrote down a detailed candidate. It appears to be the appropriate linear homotopy type theory for these tangent infinity-toposes. We have been exploring it as a possible general home for quantum programming languages.

Interviewer: Category theorists often use languages such as Agda, Lean, or Haskell rather than Python or C++. Do those languages already have the properties required here?

Urs Schreiber: We have to be careful with "required." The full linear structure is not present in the homotopy type-theoretic setting we are discussing, but plain homotopy type theory can already do something useful.

In our article "Topological Quantum Gates in Homotopy Type Theory", we describe how homotopy type theory, without the additional linear quantum structure, can encode and verify topological quantum gates.

What it cannot intrinsically certify is every specifically quantum constraint, such as the no-cloning principle. A suitable linear type theory could certify that a program respects those constraints. In plain homotopy type theory, one can also write processes that are not quantum processes, so the language alone does not provide that guarantee.

Our theorem indicates that a language such as Cubical Agda could express otherwise complicated topological quantum gates very elegantly. A further question is whether one could extend it to reflect the tangent infinity-topos structure: something like a cubical, linear, parametrised version of Agda.

I do not have an implementation of that. It is an ambitious question, but a reasonable one.

Interviewer: In your earlier interview with *The Reasoner*, you compared the excitement around the marriage of homotopy theory and type theory to the dot-com bubble. As I remember it, the point was that the underlying technologies were useful, but the initial excitement ran ahead of what they could deliver. They became useful later, rather than at the moment when the stocks were soaring.

Were people getting ahead of themselves with homotopy type theory? Is it more developed and useful now?

Urs Schreiber: I do not remember my exact wording, but that is a fair point to consider. Around 2013, when homotopy type theory became widely visible, there was enormous excitement and great expectation. Those expectations have not all been realised. The excitement has subsided somewhat, but I think the promise remains. Perhaps, as in the dot-com case, the initial enthusiasm ran ahead of development.

There were two broad promises. One was a new foundation for mathematics, native to homotopy theory and therefore, in the sense we have discussed, to gauge theory. The other was practical relevance to programming and especially certified programming.

For example, people hoped it might help us handle two very different-looking codebases that nevertheless do the same thing. Homotopy type theory might make their equivalence precise and operationally useful without requiring literal identity. That particular hope has not really materialised.

Meanwhile, proof assistants, previously a niche tool, have become much more prominent in mathematics. With AI developing, I suspect proof certificates may eventually be expected for many mathematical claims. Yet that development has largely proceeded without homotopy type theory. Kevin Buzzard, who has played a major role in promoting Lean, consciously chose a different direction. There were substantial debates about this on MathOverflow.

Interviewer: Was Buzzard the person who recently formalised Fermat's Last Theorem?

Urs Schreiber: Not him, no. I was referring to Anthropic's announcement of an AI-generated Lean formalisation of an older version of Wiles' proof of Fermat's Last Theorem. As Buzzard explains in his blog post, he has a substantial grant to work on the same general problem, but his project uses different methods and approaches. So although it might look as though he has been scooped, the projects differ. Things are developing very rapidly.

Interviewer: I was really asking which applications now give homotopy type theory a more concrete role, beyond that initial excitement.

Urs Schreiber: That brings us back to topological quantum gates.

One question that remained unresolved during the early excitement was: outside homotopy theory itself, where would programmers actually need homotopy? A language that natively understands higher gauge symmetry is interesting to mathematicians, but where is its real-world or industrial application? There were suggestions, but they did not become concrete.

Our article makes the following observation: topological quantum computing is precisely a kind of computation in which homotopical structures are essential.

We have already seen why. The gates are topological because they involve global processes in an IR-completed theory. Their description uses classifying spaces and the topology of configuration spaces. That is exactly the structure homotopy theory, and therefore homotopy type theory, understands.

Our theorem shows that monodromy quantum gates defined by Knizhnik-Zamolodchikov connections, whose classical analytic description requires a substantial sequence of constructions, become native language operations in homotopy type theory.

The key operation is monodromy: parallel transport around a loop, of the sort with which we began the conversation. Homotopy type theory does not need to be taught this operation from scratch. It is built into the notion of type transport.

Dependent types express the idea of spaces fibred over other spaces. Type transport describes how the fibre moves when we follow a path in the base. That is exactly what is needed for these topological quantum gates.

Homotopy is not a central native ingredient of ordinary classical programming or standard quantum circuit programming. But it is right at the centre of topological quantum computing. And topology may well be needed to stabilise quantum systems sufficiently for useful computation.

So I suspect-and this is a speculation, which may sound bold to other experts-that topological quantum computing could be the "killer application" of homotopy type theory as an actual programming language.

25. Type transport and certified quantum compilation · 2:09:40

Interviewer: Is the connection between type transport and physical applications a recent development?

Urs Schreiber: Parallel transport itself is an old idea. It underlies the Aharonov-Bohm effect you mentioned. Experts also know that Gauss-Manin and Knizhnik-Zamolodchikov connections are relevant to the braid transformations of nonabelian anyons.

What our article points out is that this established theory has a synthetic incarnation in homotopy type theory. Structures that have to be defined and constructed laboriously in the classical account are already present in the language.

That matters because type-theoretic programming can produce code with a certificate of correctness relative to a stated specification. If the relevant operations are native to the language, their use can be certified there without first adding an entirely separate framework.

We have not implemented this yet; the work is currently on paper. But it could have a significant impact on topological quantum compilation.

Software certification is often neglected in industry. People write code containing bugs, ship it, and fix the bugs over time. There is a long-running debate about whether one should instead prove in advance that a program satisfies its specification.

For quantum computing, there are strong reasons to expect certification to become much harder to avoid.

First, the ordinary style of debugging does not transfer directly. In a classical program, one can pause execution, inspect the registers, and check whether the machine is in the intended state. In a quantum computer, inspecting a register means measuring it, which changes the quantum state and disrupts the coherent computation. We cannot simply stop and inspect the system in the same way.

Second, compilation will be substantially more complicated, particularly for topological quantum computers. Even today, a systems programmer can sometimes inspect or write a small amount of assembly code and understand the machine's elementary operations.

For a topological quantum computer, translating a desired quantum circuit into the required dance of anyonic braids can produce something far beyond intuitive human inspection. Our article includes pictures of the braids corresponding even to simple gates. One cannot expect to look at a complicated braid pattern-the topological analogue of machine code-and immediately understand which computation it performs.

That makes the translation itself a significant source of possible errors. We need to know that the compiled motions really implement the intended circuit.

Finally, useful quantum computing time will be expensive. We cannot count on running a program thousands of times, each time discovering another avoidable error, before getting it right.

For these reasons, I expect formal verification to become essential for useful quantum computers, especially topological ones. That motivates a type-theoretic language that natively supports topological quantum gates. Our article argues that homotopy type theory provides such a foundation.

The further step is more speculative. Homotopy type theory alone does not intrinsically certify linearity or no-cloning. But its categorical semantics point to tangent infinity-toposes, and a linear version of homotopy type theory exists on paper. There are therefore good indications that an implementation-in an extension of Agda or another proof assistant-could provide a foundation in which certified quantum programming languages are embedded.

26. Writing, research, and the nLab · 2:15:52

Interviewer: I have to ask about your research process. You maintain the nLab, which contains a great deal of advanced mathematics and physics, and you write extensively. Has the writing process itself led you to breakthroughs?

Urs Schreiber: Do you mean writing on the nLab specifically, or writing in general?

Interviewer: In general.

Urs Schreiber: Absolutely. Writing is how I communicate with myself. I start writing, read what I have written, and rewrite it. It is rather like a computer writing to memory so that it can read the information back. That is how I work.

Interviewer: You began with Usenet, then the n-Category Café, and then the nLab?

Urs Schreiber: That is right. I joined those discussion forums because I lacked people nearby with whom I could discuss what I was researching. Nobody around me was working on the same things.

Interviewer: You were studying string theory at a university without other string theorists?

Urs Schreiber: Yes. My diploma thesis-what would now be called a master’s thesis-was on nonabelian strings, which is where some of these interests began. Nobody else in the department was doing closely related work. My adviser had an interest in these questions and, I think, enjoyed following the thesis as an opportunity to learn something himself.

So I was largely working on my own and looking online for people to talk to. This was the period of the free Usenet, before corporations took over much of online discussion.

As Usenet declined, blogs became prominent. But I found that blogs were not well suited to the discussion I needed. They are mostly organised around one person making an announcement, followed by comments-a largely one-way format.

Eventually, I had the idea of a shared place to record the notes we kept producing. That is how the nLab began.

27. Reading Hegel and recognising categorical structure · 2:18:40

Interviewer: Let me ask about reading, then. What gets lost when Hegel is read in translation? You have described his writing as poetic.

Urs Schreiber: Perhaps we should explain why Hegel has suddenly entered the conversation! But to answer your question, I have read him in the original German, and I would describe the writing as something akin to poetry.

It does not rhyme, of course. But despite presenting itself as scientific, it often reads more like a kind of spiritual or philosophical poetry than a scientific text. It is already difficult in German, in the way a poem can be difficult: much may lie between the lines. I imagine it must be particularly challenging to read only in translation.

Interviewer: You also mentioned that Lawvere read Hegel. Has Hegel influenced your work throughout, or is that a more recent influence?

Urs Schreiber: It happened in an unusual order. The influence was almost retrospective.

When I was discussing the beginnings of cohesive topos theory online, David Corfield, a philosopher, kept saying that what I was describing reminded him of Hegel and that I should read him.

At that point I knew almost nothing about Hegel, and I did not particularly want to know. The vague connection to Marx that I remembered from school did not appeal to me. I kept ignoring the suggestion, but David kept returning to it. Eventually I thought I would look just to put an end to the discussion.

Instead, I was struck by what I found. I went through a rather intense period of reading Hegel day and night and comparing the text with topos theory.

Lawvere, whom we discussed earlier in connection with synthetic differential geometry, had also been interested in Hegel. He regarded himself as a Marxist, although the political aspect is not what matters for this discussion. What interests me in the *Science of Logic* is its ontology.

Lawvere suggested in writing that some of Hegel’s mysterious-sounding ideas could usefully be translated into categorical constructions. In effect, though these were not his words, he proposed a way to formalise aspects of Hegel.

Hegel famously says things that seem to defy ordinary logic. Without an openness to what he may be trying to do, it can be difficult to see what he means. But there is a recurring structure of opposed or dual concepts, and Lawvere showed how a number of examples could be related to adjunctions in category theory.

I think adjunctions are the heart of category theory. The theory is not ultimately just about categories, functors, and natural transformations. Much of its substance appears at the next step, with adjunctions. An adjunction expresses a precise relationship between two complementary constructions, a kind of duality.

Lawvere took this further, considering patterns of adjunctions and relating them to patterns in Hegel. When I began reading Hegel, I knew enough topos theory to appreciate what Lawvere was saying. But I felt that he had touched only the first stage of a much longer progression in the *Science of Logic*. I began trying to match further stages with constructions in higher topos theory.

Interviewer: Is that where the idea of being arising from nothing comes in?

Urs Schreiber: That is where the progression begins.

Very roughly, one can formalise a unity of opposites through an adjunction between modalities. A modality is a structure that singles out an aspect or quality of the objects in a topos.

For example, a flat modality, when available, captures an underlying discrete aspect of a generalised space-its points, in the relevant sense. A shape modality captures its homotopy type: how it hangs together globally, as a classifying space does.

Adjunctions between modalities express a precise relation between these different qualities. Shape sees global connected structure; flatness sees a discrete aspect. An object can be understood through the relationship between these poles.

In the *Science of Logic*, Hegel develops a progression in which opposed qualities are brought into a new relationship, giving rise to another quality with a further opposite. Calling these “contradictions” can be misleading; perhaps “tensions between qualities” is better. He describes a kind of logical process through which increasingly rich structure appears.

Lawvere pointed out that mathematical structures can match aspects of that process. We can then ask what its most elementary starting point would be.

Every topos has a terminal object, which we can think of as a point, and an initial object, which we can think of as empty. There are constant functors sending every object to one or the other, and these fit into an adjunction. The point suggests pure being without further internal structure; the empty object suggests nonbeing.

In that precise categorical sense, one can read the adjunction as a formal counterpart of Hegel’s opening opposition. One can then ask whether there is a repeated progression of adjoint modalities producing further qualities. Such patterns can be related to cohesive toposes. I found that connection fascinating.

28. Categorical duality and physical duality · 2:28:45

Interviewer: Is this similar to dualities in physics, such as AdS/CFT? There, a calculation that is difficult in one theory can become tractable in a dual theory. Is categorical duality the same kind of idea?

Urs Schreiber: Some physical dualities are indeed expressed by adjunctions. I am not sure about all of them.

One example is M/IIA duality, which is related to the project I mentioned at the beginning and am currently finishing. More generally, the relationship between a higher-dimensional theory and its Kaluza-Klein reduction can be formulated through an adjunction, particularly when one works with the IR-complete theories.

Interviewer: Would you say your notion of duality is more general, with some physical dualities as special cases?

Urs Schreiber: I would not assert a general inclusion in either direction. There is at least an overlap.

The categorical notion is also close to the broader philosophical use of “duality” for complementary concepts. Even simple examples involving even and odd numbers can be phrased through adjunctions between appropriate constructions.

Interviewer: And an adjunction is the same thing as a pair of adjoint functors?

Urs Schreiber: Yes. These phrases describe aspects of the same structure. A pair of adjoint functors forms an adjunction; each functor participates as one of its two sides.

29. Mathematics as a way to see physical structure · 2:31:16

Interviewer: On another podcast, *No Time*, you described how, once axioms are written down, the possible proofs are in a sense already there, waiting to be found.

You work closely with mathematical structures that many physicists rarely encounter. Do you have a sense of how much mathematics nature uses? Does it use all mathematics? Complex numbers and tensor calculus were developed before their eventual physical applications; it can seem that nature got there first.

Or, to put it differently, what have you learnt about physics by studying mathematics that many physicists do not use?

Urs Schreiber: That is a difficult question. To my mind, mathematics is like a microscope for the internal structure of physics.

We should distinguish two senses of “mathematics.” One is calculation within a framework we already understand. We might have a Calabi-Yau manifold and want its intersection numbers. Much of the work at a string-math conference has this character: the mathematical setting is known, and one seeks more examples or more detailed calculations within it.

The other question comes before calculation: what mathematical language expresses the physical idea in the first place? This is the kind of question Einstein faced in bringing differential geometry into physics, or the founders of quantum mechanics faced in finding the right framework for quantum states and observables, including von Neumann’s use of Hilbert spaces.

Before solving an equation or calculating an example, we may need to discover what the mathematical theory should be. That is primarily the sense in which I am using “mathematics” here.

At the beginning of the twentieth century, there were fortunate coincidences. When Einstein needed differential geometry, it had already been developed and could be explained to him. When quantum theory needed Hilbert spaces, they too were available from earlier work undertaken for somewhat different reasons.

Those coincidences enabled enormous progress. I suspect that their absence in other areas partly explains the sense of stagnation in the later twentieth century. Physics continued to develop, but the appropriate mathematical language did not always keep pace.

Part of my claim is that structures such as higher toposes are missing from the usual toolbox, and that this lack contributes to our difficulty with nonperturbative phenomena.

30. Is string theory too mathematical-or not mathematical enough? · 2:35:33

Interviewer: People sometimes say the Standard Model works too well. They also say string theory is too mathematical. Yet you argue that we do not have enough mathematics even to define nonperturbative physics properly, although nature itself is nonperturbative. We need better definitions.

Urs Schreiber: I think different meanings of “mathematical” are being confused.

When people say string theory is too mathematical, I think they often mean that it is too detached from physical applications: too much speculative theoretical physics. But speculation, or even formal-looking speculation, is not automatically mathematics.

For a mathematician, mathematics involves clear axioms and rigorous deduction. String theory certainly does not suffer from an excess of those; in important respects it lacks them.

Both criticisms can therefore be true. As a physical theory, string theory has not been connected sufficiently to actual physics. As a structure in mathematical physics, it is not yet mathematical enough.

People respond by mentioning mirror symmetry, derived categories of D-branes, and so on. Certainly, string theory has stimulated important mathematical developments by suggesting structures and relationships worth studying. But the resulting pure mathematics does not necessarily clarify string theory itself. Work on homological mirror symmetry, for example, can become an independent mathematical subject with a string-theoretic origin, while leaving fundamental questions about the physical theory unresolved.

The major difficulty since the so-called second string revolution has been the lack of foundations for the new objects being discussed.

The first revolution rested, at least in principle, on a reasonably understood framework: worldsheet perturbation theory, superconformal field theories, their correlators, and the extraction of low-energy effective theories, including supergravity theories and anomaly cancellation.

Around 1995, attention shifted to D-branes and M-theory. People began discussing branes as though their nature were already established. The discussion was no longer grounded entirely in worldsheet analysis, but a comparably complete replacement formalism was not supplied.

There are proposals for a Dirac-Born-Infeld action for coincident D-branes, for example, but these do not straightforwardly do everything required. Nevertheless, discussions often proceed as though coincident branes were already completely defined.

The test is whether one can hand the definition to a mathematician and say, “Here is the object; work out its properties.” For much of this structure, that definition has been missing. In that sense, the second revolution-branes and nonperturbative dualities-still awaits its mathematical formulation.

31. What would a theory of everything explain? · 2:39:24

Interviewer: For my last question: some people say we do not need a theory of everything. Do we?

Urs Schreiber: In one sense, that is what science is about. We construct theories of more and more phenomena, and the limiting aspiration is to understand everything we can.

But one could also argue from the opposite direction that the physicists’ “everything” is not extensive enough. Usually, a theory of everything means unifying gravity with the other three fundamental forces.

We can ask questions further down the ontological ladder. Before asking how to unify forces, why do we have fields at all? Why spacetime? Why mathematics? Why should things be described by mathematical structures in the first place?

Interviewer: What do you say to yourself when thinking about those questions? Why mathematics? Does spacetime emerge? Is it discrete?

Urs Schreiber: Those are distinct questions. Let me tell a story about the mathematical one, which will, unsurprisingly, lead back to higher topos theory.

When homotopy type theory became prominent, one striking aspect was that it made something simpler. We have encountered that repeatedly in this conversation: a complicated and unwieldy construction can become natural once we change perspective.

Martin-Löf had developed his type theory from very elementary logical considerations. In his lectures on the origins of logical operations, he gives a compelling account of the basic structures needed to reason at all. To speak about “this and that,” for instance, we need a way to combine the relevant data. The theory can seem almost like a logical minimum.

But an issue appeared at a particularly interesting point: identity. In type theory, things have types, so identities between things should also have a type. Given two terms, a proof or witness identifying them is itself data, belonging to their identity type.

One might expect that whenever two terms are identified, there is essentially only one witness of their identity. Yet that uniqueness did not follow from the basic rules. Adding it as an extra principle caused other difficulties. For years, the status of identity types remained a source of tension.

Homotopy type theory arose, in part, by stepping back and accepting what the formalism was telling us. Why insist that witnesses of identity be unique?

In the physical language we have been using, an identification is like a gauge transformation. Two configurations may be related in more than one way. The transformations themselves can then be related by higher transformations. Identity types can have identity types, and so on.

From this perspective, the apparently troublesome feature is exactly what we want. Higher gauge structure was present in the formalism all along. We needed to recognise it rather than suppress it. The univalence axiom adds another essential ingredient, but much of the conceptual shift is simply to accept a structure that is richer than our initial expectations.

That is a fascinating story when we ask where mathematics and physics originate. Starting with very basic logical operations-things, types of things, pairs of things, identities-we find ourselves led to a setting naturally described by infinity-toposes, with higher gauge symmetries already built in.

One can then pursue the Hegel-inspired question. An infinity-topos has the initial opposition between the empty and terminal objects. Can one follow a progression of modalities corresponding to discreteness, continuity, and other aspects of geometry?

What is striking is that this concerns the emergence of the concept of space, before the emergence of any particular physical spacetime. Logical structure can lead us to the notion that there are spaces at all. Which space might describe our world is then a further question.

Interviewer: So space and time emerge from logic?

Urs Schreiber: So far, I was speaking about space in the mathematical sense, rather than physical spacetime. For spacetime, I have another story: emergence from the superpoint. Have you seen that?

Interviewer: Yes.

Urs Schreiber: Very briefly, we begin in an infinity-topos and investigate that progression of modalities. This is part of what my work on differential cohomology in a cohesive infinity-topos concerns. At the final stage of the progression under consideration, supergeometry appears.

A superpoint is pointlike in the ordinary geometric sense but has odd infinitesimal directions, of the kind used to describe fermions.

We can then perform a further mathematical construction: take universal invariant higher central extensions. This produces a tower of objects growing out of the superpoint, somewhat like a Whitehead tower.

I am referring here to the article “M-Theory from the Superpoint,” which I wrote with John. It makes this construction precise. The superpoint plays the role of an initial seed, almost a cosmic egg, and the extension process reveals increasingly rich structure within it.

More precisely, beginning with the N = 2 superpoint, the first universal invariant central extension in the sense of that article gives three-dimensional super-Minkowski spacetime. At that stage, it is reasonable to speak of actual spacetime structure: a Lorentzian metric appears, as the article explains.

We then repeat the process. Starting with N = 1 supersymmetry in three dimensions, we double the fermionic directions and take the next extension, obtaining N = 1 super-Minkowski spacetime in four dimensions.

The progression continues through the critical spacetime dimensions associated with the Green-Schwarz superstring and ends in eleven dimensions. It also reveals the M2- and M5-brane structures and their relation to the rational homotopy type of the 4-sphere.

That is the sense in which I would say spacetime geometry emerges in this construction. It may sound extraordinary when stated briefly, but I invite people to read “M-Theory from the Superpoint.” The claims there are formulated as precise mathematical statements.

32. A reading suggestion · 2:50:48

Interviewer: Thank you so much for your time. Do you have any book recommendations-philosophy, fiction, novels, or textbooks? Something recent or something you read long ago?

Urs Schreiber: That is a difficult question. Anyone interested in ontology should try reading Hegel’s *Science of Logic*, starting on page one.

Interviewer: That is already one recommendation. Good!

33. Where to find Urs · 2:51:22

Interviewer: Where can listeners find you?

Urs Schreiber: I have a personal page on the nLab. Searching for my name and "nLab" should bring it up, and it contains further links.

I post on X, formerly Twitter, and sometimes on YouTube. It is perhaps too much to call it an active channel, but I often make preliminary recordings of talks, and occasionally record other people's talks. You can also find recordings of my talks elsewhere on YouTube.

I contribute to MathOverflow and Physics Stack Exchange from time to time, although less often these days. There is also the nForum, where we discuss additions to the nLab. It has been fairly quiet lately, but I help host it.

And of course people can email me-although I am busy, so please do not send too many emails!

Interviewer: Thank you very much for your time.

Urs Schreiber: Thank you.