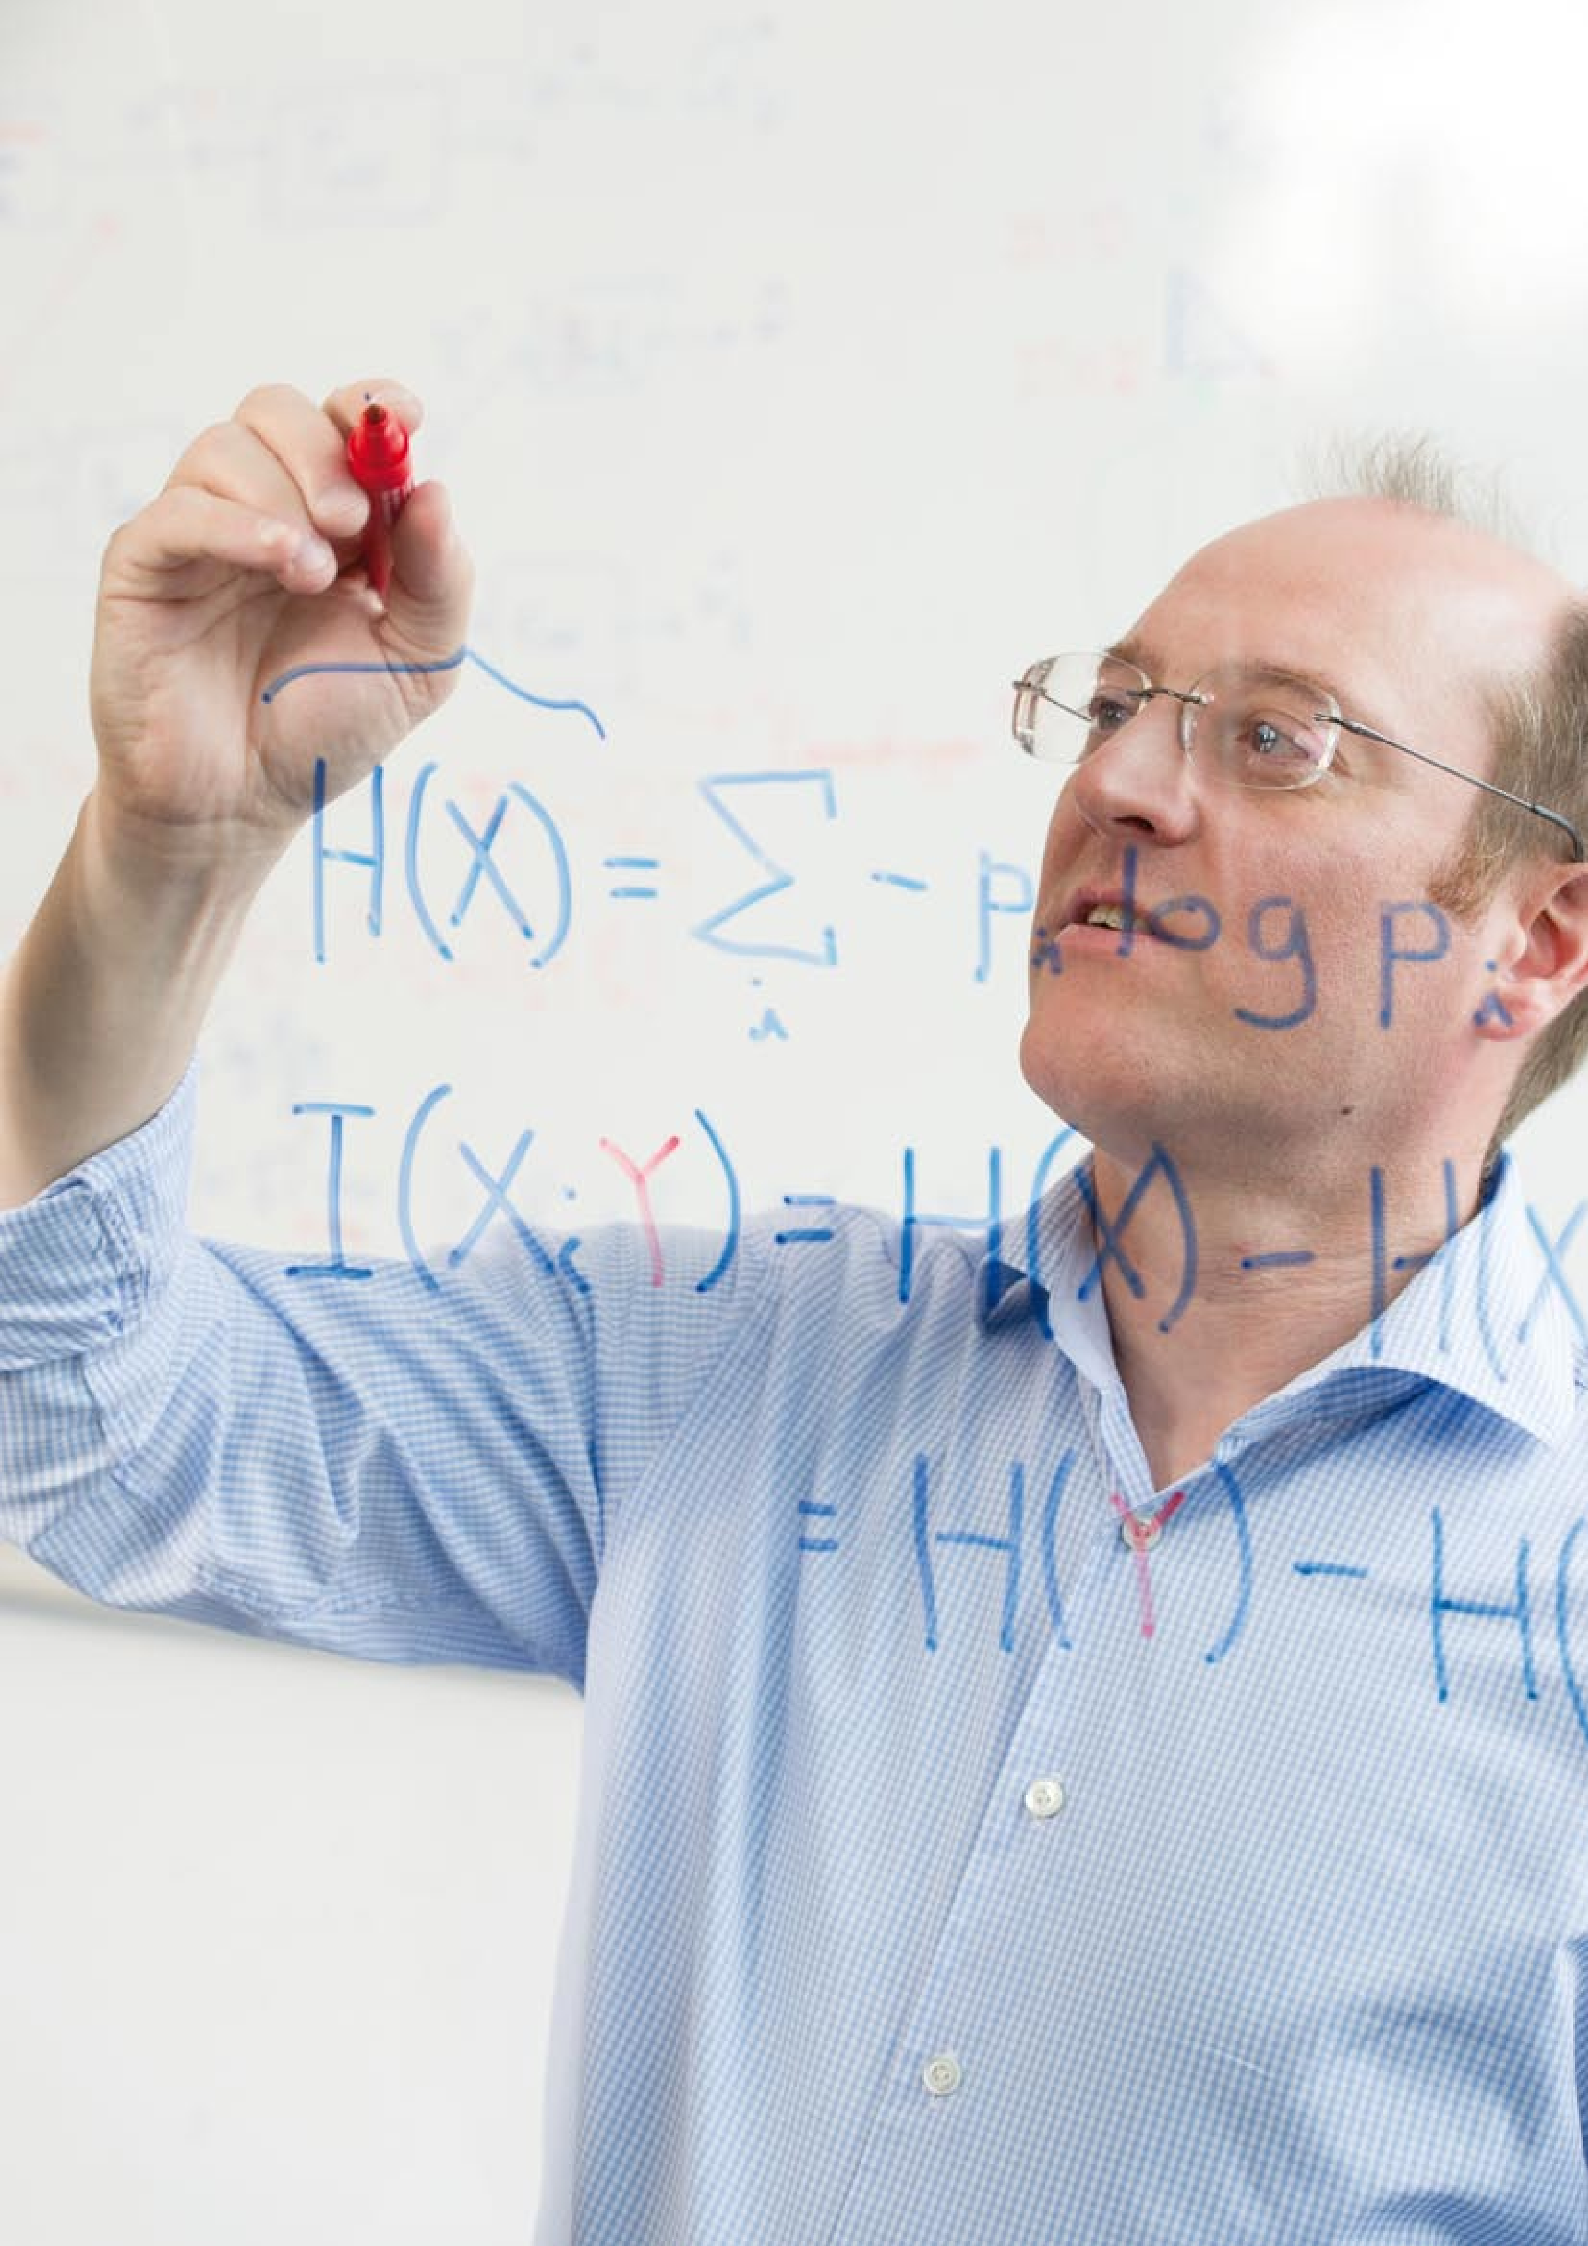

$H(X) = \sum_i - p_i \log p_i$
$I(X;Y) = H(X) - H(X$
$= H(Y) - H($

Frank Kschischang, Mansoor I. Yousefi, Gerhard Kramer, Patrick S. Regan, René-Jean Essiambre and Luca Barletta

Laser light, glass fibers, digital data: Probing the limits

Based on an interview by Patrick S. Regan (PSR) with TUM Prof. Gerhard Kramer (GK); Hans Fischer Senior Fellow Prof. Frank Kschischang, University of Toronto (FK); Rudolf Diesel Industry Fellow Dr. René-Jean Essiambre, Alcatel-Lucent Bell Labs (RJE); Dr. Luca Barletta (LB) and Dr. Mansoor I. Yousefi (MIY), both TUM.

PSR: Optical communications technology is fairly mature now, isn't it? What kinds of open scientific questions are being addressed by this Focus Group?

GK: The technology is mature in many ways, although of course the companies that provide commercial systems – basically, the extremely high-capacity fiber-optic systems that serve as the backbone of the Internet – continue to push the limits. The thing is, the *theoretical* limits are, strictly speaking, unknown. And that has been one goal of our Focus Group's research, to investigate the theoretical capacity of communication over optical fiber in a rigorous way.

PSR: The basic framework for that, information theory, was created nearly 70 years ago by a paper Claude Shannon published in the Bell System Technical Journal: "A Mathematical Theory of Communication." To what extent is this research rooted in Shannon's original work, and to what extent does it branch off in new directions?

FK: Shannon really looms over the field, and everything it touches. I remember during an ISIT meeting, the IEEE International Symposium on Information Theory, on the 50th anniversary of that paper, a giant picture of Shannon was hanging on the stage, as if he was watching you giving your talk. One guy said it was too intimidating, and he turned it around. If people only knew – you could argue that Claude Shannon has had a bigger impact on the digitalization of life today than Steve Jobs or Bill Gates.

GK: Shannon's theory still forms the basis for defining a communication channel mathematically. But optical fiber is different from other transmission media, and the fiber channel requires a special approach.

MIY: We just hosted the first Munich Workshop on Information Theory of Optical Fiber, an international conference sponsored by the TUM-IAS. The main theme of the workshop was finding a fundamental limit of the optical fiber channel. So far we have estimates and we have lower bounds and so on, but still, several decades after the introduction of the optical fiber, no one yet knows the fundamental limit on information rate in this channel.The question is still open, and it is very important. It's the cornerstone of the world's communication networks.

RJE: Nonlinear effects need to be taken into account. Optical fiber is a very transparent medium, an extremely efficient waveguide made from ultra-pure silica glass. Modern fibers attenuate light by a factor of 2 every 15 kilometers. At about every one hundred kilometers, you need to insert an optical amplifier to bring the power back up. You can transmit light for hundreds or thousands of kilometers that way. Very small in diameter, 125 microns, it's comparable to the thickness of a human hair. It has a small 10-micron core that keeps light inside the fiber. And the light is so concentrated inside the core, there's such a high intensity there, that it changes the index of refraction of the material, for all the different users – there could be typically one hundred different "users" …

Gerhard Kramer

PSR: Meaning many different wavelengths of light carrying data simultaneously through a single fiber?

RJE: Different wavelengths, exactly, different frequencies. And because they all share the medium, with their own data, they start to interfere with each other, even if they don't overlap spectrally. Your distortion depends on the other user's data. And that creates distortion for you. And that's part of the reason for the Focus Group, to discover how much information can be sent in the presence of these distortions created by nonlinear effects.

PSR: The TUM-IAS framework has enabled you to bring together people with different expertise but common interests, younger as well as more senior, from industry as well as from universities. Where are the various Focus Group members "coming from" scientifically?

GK: Frank and I are information theorists, with somewhat overlapping research programs. For us at TUM, we have a large group, especially through the funding that's been available through my Humboldt professorship. We have active research on a very wide range of topics, mainly in reliable communications, secure communications, and data compression.

RJE: My background is in physics, more specifically nonlinear dynamics, chaos theory, and astrophysics as well. I have worked on solitons in the past, and nonlinear switching. At Bell Labs I also work, for instance, on more applied problems of optimizing the transmission of information over our commercial systems.

LB: I did my PhD in wireless communication. I started to study these problems during an internship with René at Bell Labs. Here, with the Focus Group, I'm studying certain kinds of disturbance in optical channels introduced by lasers. These components are not ideal, so they introduce noise, and it's a particular kind of noise. It's a noise with memory, a multiplicative noise, that is different from the usual additive noise. Phase noise arises not only from the lasers, but also because of the nonlinear effects in the glass fiber. Now, the main goal in communication is to detect the data that was transmitted, and so you want to compensate for noise that disturbs the communication. But to do this – or potentially even to use the noise to carry information – we need to understand the noise better than we do now. That's my focus.

MIY: I'm an engineer but with emphasis on the mathematical side. I did my PhD with Frank at the University of Toronto in Canada. Initially I intended to work on probabilistic and graphical models; I had no background in optical communications. For a year and a half I worked on network coding and related topics, before Frank suggested that I look at the fiber problem. I got attracted by the rich mathematics of the nonlinear dispersive partial differential equations that appear in this context. We also learned the nonlinear Fourier transform, which turns out to be an elegant and powerful tool for dealing with wave propagation in certain nonlinear media such as optical fiber. Then we started to see how we could actually use it to intelligently communicate over what's called the nonlinear Schrödinger channel.

PSR: How did this collaboration come together?

GK: The idea for the Focus Group started with Frank. Just when I was about to join TUM, Frank approached me at a conference and told me he had a sabbatical coming up. That was a natural fit. I was looking for topics at the time, and I felt it would be an opportunity to bring together an industry Fellowship and an academic Fellowship. If I put the two together, it was just natural to choose this topic, because it's very strongly industry-related. I had been working with René at Bell Labs, and Frank, I knew, was working on the fiber side and more, with a bit of a different spin, the coding side, coding from various perspectives. He had visited Bell Labs while I was there, and he had interacted with René already. It was a natural fit, and through the TUM-IAS it was very easy.


Mansoor I. Yousefi

PSR: I gather you see this as the start of something bigger. Why now?

GK: It hasn't been studied much in the information theory community, which is really unfortunate, because the problems are challenging, they're very practical, and you have really good people who could contribute.

FK: There's a really good reason that we, information theory people, should come in now. And that's because the physics has matured to where the channel model is more or less agreed upon. The devices haven't changed, basically, the fiber itself hasn't changed in the last decade or so, but there

Frank Kschischang

have been improvements in the physical understanding of the channel. Now you need information theoretic progress to finish the story.

MIY: That's right. And again, importantly, we have this limited capacity at the core of the Internet, facing ever increasing, virtually limitless demand. That makes the information theory of optical fiber very interesting.

GK: Information theory, coding in particular, always comes in when you need to improve efficiency. Up to a point you can use much simpler methods, but once you start hitting these boundaries, you can *shift* the boundaries through more sophisticated methods. And this shifting is quite strong. It's always a good step beyond what you could do without it.

RJE: And the problem itself is actually very new to the information theory group, isn't it?

GK: It is.

RJE: Because there's something that you never dealt with. You cannot do a simple extension of what people have done with other media. It's very difficult to do. So at the fundamental level, it is also something that is open for a lot of discoveries and new understanding.

PSR: When Frank says that the physical model has improved, what does that mean exactly? Is it correct to say that what you mean by a model is a set of interlocking equations describing separate physical phenomena that interact?

MIY: A set of mathematical equations that have predictive capabilities. To predict what's going on, precisely enough to be useful in practice.

RJE: There are a lot of things happening in the fiber. The model we are using, it's a very good model, but it's still not exact. There are other phenomena – not only photons, there are also phonons, mechanical waves, and they interact and lead to some other effects that we neglect here – the question is, *can* we neglect them? And so all this goes into refining the model for our purposes. In fibers, as soon as you go to a certain level of power, it could be only five or six dB higher, something else is happening, and the model now starts to fail. So you actually have to know how the model changes, according to where you're pushing the limits.

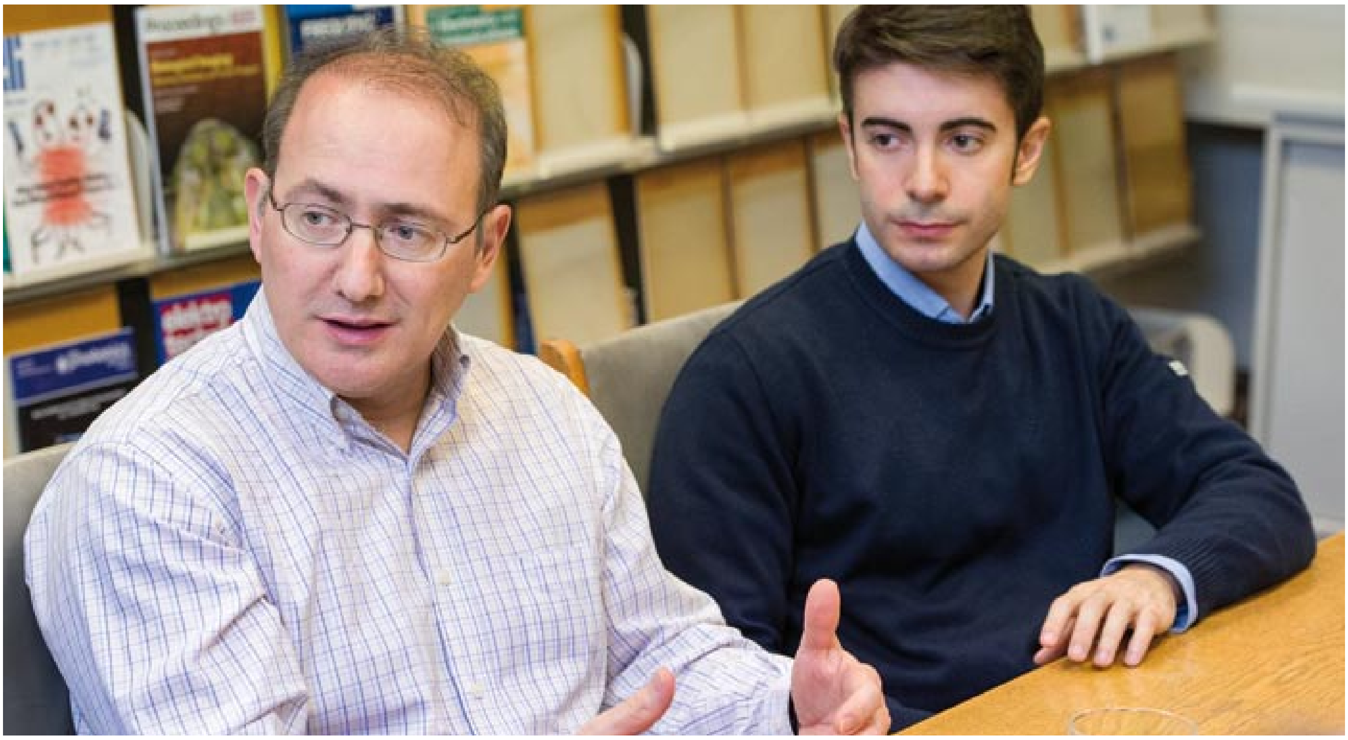

René-Jean Essiambre and Luca Barletta

PSR: And are you looking also at optical amplifiers, components like that?

RJE: Absolutely. I'm not the most expert in that, but I do have some expertise in terms of all the noise that is generated. There are others, more expert than I am, that I can talk to. That way we bring the knowledge together, and then we figure out what is the proper model and what the limitations of this model are.

FK: That's right. We need the physicists to give us the model to some extent. Our starting point is the mathematical model, and we take off with the mathematical approach from that point. But we need people like René to tell us what the model really is and to delineate the applicability of the models, so we don't go and spend two decades solving the wrong model.

PSR: In some sense this whole pursuit takes off from a practical concern – a "capacity crunch" in the face of growing Internet traffic. How do you think this theoretical research might connect with technological applications in the future?

LB: I think a step could be, at least from my side, getting a better understanding of these channels with memory, and especially devising algorithms, practical algorithms for achieving capacity limits. It will be very important to devise low-complexity algorithms, because right now what theory suggests is to use very complex algorithms. So what happens if we use low-complexity algorithms? Will we still come close to the limits, or not? This is a good question to answer.

FK: This elegant nonlinear Fourier transform theory that Mansoor has developed not only helps us understand the limits, but also could suggest ways to achieve capacity closer to the limits.

MIY: Our work would predict that you can achieve these high data rates in the fiber if you use these particular methods. Looking forward, one direction is to continue the approach that we started. But we are looking at other ways as well. There's not just one way you can look at research problems.

PSR: Do the answers you're seeking have implications for system design, device design, economics of networking?

RJE: The optical hardware already is quite mature. We know how to do low-noise amplification. Coherent detection works very well.

GK: But maybe, I hope, the math will influence the design of some future devices. In relation to dispersion management, for example, it already has. So, for instance, you have to wonder if there could be new kinds of filters – rather than being wavelength filters designed for sinusoids – that might enable a so-called ROADM, a reconfigurable optical add-drop multiplexer, based on the math of the nonlinear Fourier transform.

FK: That's the number one burning question for us.

MIY: That is, whether or not the mathematics could enable you to do optical processing, for network routing, in the nonlinear Fourier domain.

GK: There are different classes of equations, and every one of them has some applications. Whether for communications or not, you don't always know. But you get different solutions for different classes of equations.

RJE: And there are different materials. The fiber has a nonlinearity of a certain type, so-called Kerr linearity, because the glass is a disordered medium, an amorphous medium. There are media that are different. A medium might be found that is low-loss but has nonlinearity of a different kind. We don't know. In the future this may happen.

FK: And all this goes out the window.

GK: No!

RJE: We never know. We don't see it now, but we never know.