UNIVERSITY OF SOUTHAMPTON

FACULTY OF PHYSICAL AND APPLIED SCIENCES

Optoelectronics Research Centre

TRANSFER OF OPTICAL FREQUENCY COMBS OVER OPTICAL FIBRE LINKS

by

Giuseppe Marra

Thesis for the degree of Doctor of Philosophy

February 2013

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF PHYSICAL AND APPLIED SCIENCES
Optoelectronics Research Centre

Doctor of Philosophy

TRANSFER OF OPTICAL FREQUENCY COMBS OVER OPTICAL FIBRE LINKS

by Giuseppe Marra

In just over a decade the optical frequency comb technique has completely transformed the field of frequency metrology. These devices have made the measurement of the frequency of light a much easier and affordable task when compared to with earlier techniques. With both research and technology development on these devices becoming more mature, optical frequency combs have been affecting other science areas. Applications are already found in spectroscopy, attosecond physics and astrophysics and more science and engineering areas can be expected to be affected in the near future.

The dissemination over optical fibre of optical frequency combs between research labs, or between research labs and industry, could facilitate and accelerate this process. In particular, with optical frequency standards currently exhibiting a fractional accuracy better than 10^{-17} and optical frequency combs making this accuracy available across a wide spectrum, new experiments could be devised in a wide range of research fields if ultra accurate microwave and optical frequencies were to be made available beyond the walls of metrology laboratories.

However, before the work presented in this thesis, limited research was undertaken to test how accurately an optical frequency comb could be transferred over optical fibre. Enviromentally-induced noise in the fibre, dispersion issues and other processes taking place during the propagation and detection of the optical signal could all degrade its quality to a level incompatible with the desidered applications. The experiments described here demonstrate that optical frequency combs can be disseminated over optical links, from several-km to many tens of km-long, whilst preserving the stability and accuracy of its mode frequency spacing and mode frequency to a level compatible with the majority of the most demanding frequency metrology applications.

To my Dad, who introduced me to the beauty of science.
And to my Mum, whom science did not manage to save.

Contents

List of Figures

List of Tables

Declaration of Authorship

I, Giuseppe Marra , declare that the thesis entitled *TRANSFER OF OPTICAL FREQUENCY COMBS OVER OPTICAL FIBRE LINKS* and the work presented in the thesis are both my own, and have been generated by me as the result of my own original research. I confirm that:

- this work was done wholly or mainly while in candidature for a research degree at this University;
- where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
- where I have consulted the published work of others, this is always clearly attributed;
- where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
- I have acknowledged all main sources of help;
- where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
- parts of this work have been published as: (Marra et al., 2010), (Marra et al., 2011) and (Marra et al., 2012)

Signed:.....

Date:.....

Acknowledgements

Hidden behind the plots and diagrams of this PhD thesis there are a great deal of enthusiastic and passionate discussions, long hours spent in the lab, frustrating failures, exciting ideas that turned to be wrong and, fortunately, a few that did work and led to some memorable Eureka moments. I would like to thank those people that shared all of this with me.

Firstly, I would like to thank my supervisor David Richardson, who has supported me for the whole duration of my PhD and always found time for me despite his incredibly busy schedule. I would also like to thank Periklis Petropoulos who has been really helpful expecially in the first stages of my PhD.

A major thank goes to two special people: my NPL colleague Helen Margolis who has given me great help throughout the PhD and with whom I always enjoyed having very informative discussions and sharing my enthusiasm and passion for science; and my ORC colleague Radan Slavik in whom I found not only an incredibly supportive scientist during my experiments in Southampton (and thereafter), but also a good friend with original thoughts about life.

I would also like to thank Stephen Lea for initiating work on frequency transfer over fibre at NPL, Ben Parker, Luke Jonhson and Maurice Lessing who made our lab a fun place place to be in and all the other NPL and ORC staff I have interacted with.

Finally, I would like to thank Patrick Gill who does a lot of work behind the scenes to secure the funding that enables us to enjoy the pleasures of science.

Nomenclature

AOM	Acousto-Optic Modulator
$CERN$	Conseil Européen pour la Recherche Nucléaire
CEO	Carrier-Envelope Offset
CW	Continuous Wave
CIR	Circulator
DCF	Dispersion Compensating Fibre
DSN	deep Space Network
DDS	Direct Digital Synthesizer
$EMRP$	European Metrology Research Programme
$EDFA$	Erbium Doped Fibre Amplifier
FFT	Fast Fourier Transform
GPS	Global Positioning System
$LNGS$	Laboratori Nazionali del Gran Sasso
LPL	Laboratoire de Physique de Lasers
$JILA$	Joint Institute for Laboratory Astrophysics
JPL	Jet Propulsion Laboratory
NF	Noise Figure
$NIST$	National Institute of Standards and Technology
NMI	National Measurement Institute
$NMI(AU)$	National Measurement Institute (Australia)
$NMIJ$	National Metrology Institute of Japan
NPL	National Physical Laboratory
ORC	Optoelectronics Research Centre
PSD	Power Spectral Density
PTB	Physikalisch-Technische Bundesanstalt
SMF	Single mode Fibre
SNR	Signal-to-Noise Ratio
SSB	Single Side Band
$SYRTE$	Systèmes de Référence Temps-Espace
UCL	Unversity College London
ULE	Ultra-Low Expansion
UWA	University of Western Australia

Chapter 1

Introduction

On September 2011 a press release by a research group at the Laboratori Nazionali del Gran Sasso (LNGS) in central Italy and the European Organization for Nuclear Research (CERN) in Switzerland attracted the attention of the whole scientific community. They announced that neutrinos shot from CERN always arrived at the detector in the underground laboratory in Italy some sixty nanoseconds faster than a beam of light would take (Adam et al., 2011). The early arrival implied that the neutrinos travelled faster than the speed of light, a violation of Einstein's theory of special relativity, a cornerstone of modern physics. The news spread around the world within moments (almost faster than the speed of light one might joke) receiving very wide coverage from within the media as, if the results had been confirmed, the laws of physics as we currently know them would need to be rewritten. Quite embarassingly, months after the news outbreak, two possible sources of trivial experimental errors were identified (CERN, 2012). In March 2012, another experiment, also at the LNGS, measured the speed of neutrinos to match exactly that of light (Antonello et al., 2012), to the joy of those who never believed in the early results and the disappointment of those who had hoped the experiment was the beginning of a new era of physics.

Opening my thesis with a flawed experiment might be considered an odd choice. However, I believe that, despite its flawed results, the neutrino experiment has been a gift to the time and frequency metrologist. Before September 2011, the best example of precise atomic timekeeping applications that a time and frequency metrologist could give to a non-expert audience was the Global Positioning System (GPS). Even in that case though, how precise timing leads to precise positioning has always been a difficult and laborious concept to explain. By contrast, anyone can understand the need to measure time accurately in the CERN/LNGS experiment, as this can be linked to real life experiences such as, for example, the way the average speed of a car is measured between two locations by automatic number plate recognition systems. The general public can now understand that measuring time very precisely can make the difference between a revolutionary and a flawed outcome of a fundamental physics experiment such as the

OPERA experiment. Whatever the outcome, the neutrino experiment added a powerful tool to answer the most frequently asked question to the time and frequency metrologist of why we need ultra precise clocks?[1]

1.1 Why do we need atomic clocks?

"The clock would lose or gain less than a second in some 138 million years"[2]. This is an example of how the generic press usually reports on research laboratories breaking new accuracy limits with their atomic clocks. Although this is often disliked by the time and frequency metrologists who prefer to express the accuracy in the more scientific notation of fractional parts (i.e. $a \times 10^{-n}$), the approach adopted by the press allows a wider audience to get a feel for the outstanding precision delivered nowadays by atomic clocks. After the initial "wow" effect, the question that arises in the member of the public's mind is usually "what is the point of measuring time so precisely?". Some also question whether governments should pour taxpayer money into such an apparently futile activity. In the following few paragraphs I will try to explain why precise timekeeping is becoming a more and more crucial part of our society's everyday life and how it might open whole new chapters in our understanding of how nature works. From a more individual point of view, the reason why I myself embarked in the time and frequency metrology field is simply that I find it extremely fascinating.

GPS Undoubtedly the application of atomic clocks to navigation has been the one with the biggest global impact. Nowadays everybody is familiar with this technology and a large number of us probably have a GPS receiver embedded in our smartphone. However, probably few of us are aware that the positioning technology is enabled by precise timekeeping devices on board each of the satellites. In each of the satellites a timing signal from a rubidium clock is sent to Earth and, when combined with timing signals from other satellites, allows for the position of the user to be defined. However, the accuracy of the clocks on board each of the GPS satellites is many orders of magnitude worse than that achievable today with state-of-the-art atomic clocks. With improved clocks the positioning accuracy could potentially be improved by orders of magnitude[3] allowing, for example, the possibility to land a plane using GPS signals. The applications of improved positioning systems could be as many as difficult to predict. Certainly Louis

[1] The results of the OPERA experiment could have also been confirmed or denied more quickly and more confidently if an alternative way of measuring time and frequency to GPS had been available. Of course one possible solution for an independent validation of the results could have been achieved using time and frequency transfer techniques over optical fibre links, which closely relates to the content of this thesis. To my knowledge, this solution was being taken into consideration before the latest results made it no longer necessary.

[2] News on the BBC website reporting on advances in the NPL caesium clock's accuracy - 26 August 2011

[3] Improvements in satellite-to-ground time transfer techiques will also be required in order to fully benefit from improved clocks on board of the satellites

Essen, the inventor of the first atomic clock, could not have dreamed that his invention would one day help us find a nearby restaurant using our mobile phone.

Fundamental Physics Frequency (and thus time) is the physical quantity that can be measured with the highest precision. When other quantities are related to frequency in a known way, the precision of the frequency measurement can be transferred to the other quantities. This can be used to test, for example, whether fundamental constants are really constant (Prestage et al., 1995; Dzuba et al., 1999; Dzuba and Flambaum, 2000; Lea, 2008). If this was found not to be the case, this would be a violation of Einstein's equivalence principle and could possibly provide some hints towards unifying quantum theory and the theory of relativity. Recently much attention has been put on testing the presence of time-variations of the fine structure constant α by comparing a microwave and an optical clock (Bize et al., 2002; Fischer et al., 2004; Peik et al., 2004) or two optical clocks (Rosenband et al., 2008) over year-long time scales. As optical clocks improve, detecting small variations of the fine structure constant becomes an easier task and interesting results should certainly be expected in the next few years.

Geodesy In order to account for systematic uncertainty in atomic clocks, their altitude relative to the geoid (the Earth's equipotential surface closest to the mean sea level) must be known very accurately. However, with optical clocks now reaching a relative accuracy better than 10^{-17} (Chou et al., 2010), an altitude difference of less than 10 cm can lead to systematic shifts larger than the clock's accuracy. The accuracy to which the geoid has to be known is higher than that available today. The optical clocks could therefore be used to develop a gravitational field "sensor" of much improved precision compared to any other currently existing methods (Tapley et al., 2004; Kleppner, 2007; Svehla, 2008).

1.2 A brief history of timekeeping

When did we start measuring time? Before the invention of man-made timekeeping devices, the only available "clocks" was the night-day cycle and the change of seasons. A more accurate way of telling time was introduced around 3500 B.C. by the Egyptians with the invention of sun clocks. Around 2200 B.C. the Egyptians started using the position of the stars in the sky in order to be able to also tell the time at night. Water clocks and candle clocks were introduced later, making the measurement of time an easier and more accessible task for a larger portion of society. But it was only with the invention of mechanical clocks that the art of timekeeping really took off and was sustained through the second millennium. One of the greatest contributions came at the end of the XVIII century from the English horologist John Harrison who spent a good

part of his life improving mechanical clocks, achieving the remarkable accuracy of a few seconds per day as early as 1773. The defining moments of the history of timekeeping up to 1924 are outlined in Figure 1.1.

1.3 Modern era of the art of timekeeping

The era of modern timekeeping, based on the interaction between electromagnetic waves and atoms rather than on mechanical oscillations, began in 1949 when the American scientist Norman Ramsey invented the separated oscillatory field technique. Ramsey's invention, for which he received the Nobel Prize in 1989, dramatically accelerated the rate of improvements in the accuracy of timekeeping devices. In 1955 Louis Essen demonstrated at the National Physical Laboratory in Teddington (UK) the first caesium atomic clock with accuracy superior to any other timekeeping technique (Essen and Parry, 1955). Essen's pioneering work paved the way to the development of today's ultra-accurate microwave clocks. The invention of the optical frequency comb technique in 1999 was another defining moment for the history of timekeeping of similar magnitude to that brought about by Ramsey in 1949. The invention sped up the development of optical clocks as the measurement of optical frequencies became a relatively simple process, particularly when compared to the very complex earlier frequency-chain techniques. The defining moments of this exciting modern era of timekeeping are shown in Figure 1.2.

1.3.1 Improvement of the accuracy of atomic clocks over time

Essen's first caesium atomic clock in 1955 was not the first one to use atomic transitions for timekeeping (Lyons, 1952), but it was the first to measure time more accurately than other systems based on the rotation of the Earth (Markowitz et al., 1958). His achievement had at least two major consequences. The first was that a new definition of the second became necessary. This was agreed in 1967 when the definition of the second became related to a specific transition of the caesium atom rather than a fraction of the tropical year. The second was that, still today and for many years to come, a great number of scientists would try to develop progressively more and more accurate clocks based on atomic transitions.

Today, the best atomic clocks can be as accurate as many parts in 10^{18} (Chou et al., 2010), representing an improvement of more than 7 orders of magnitude from Essen's clock and more than 8 from the accuracy achievable with systems based on the Earth's rotation (Markowitz et al., 1958). It is worth noticing that such a dramatic accuracy improvement has been achieved in just under 60 years of research.

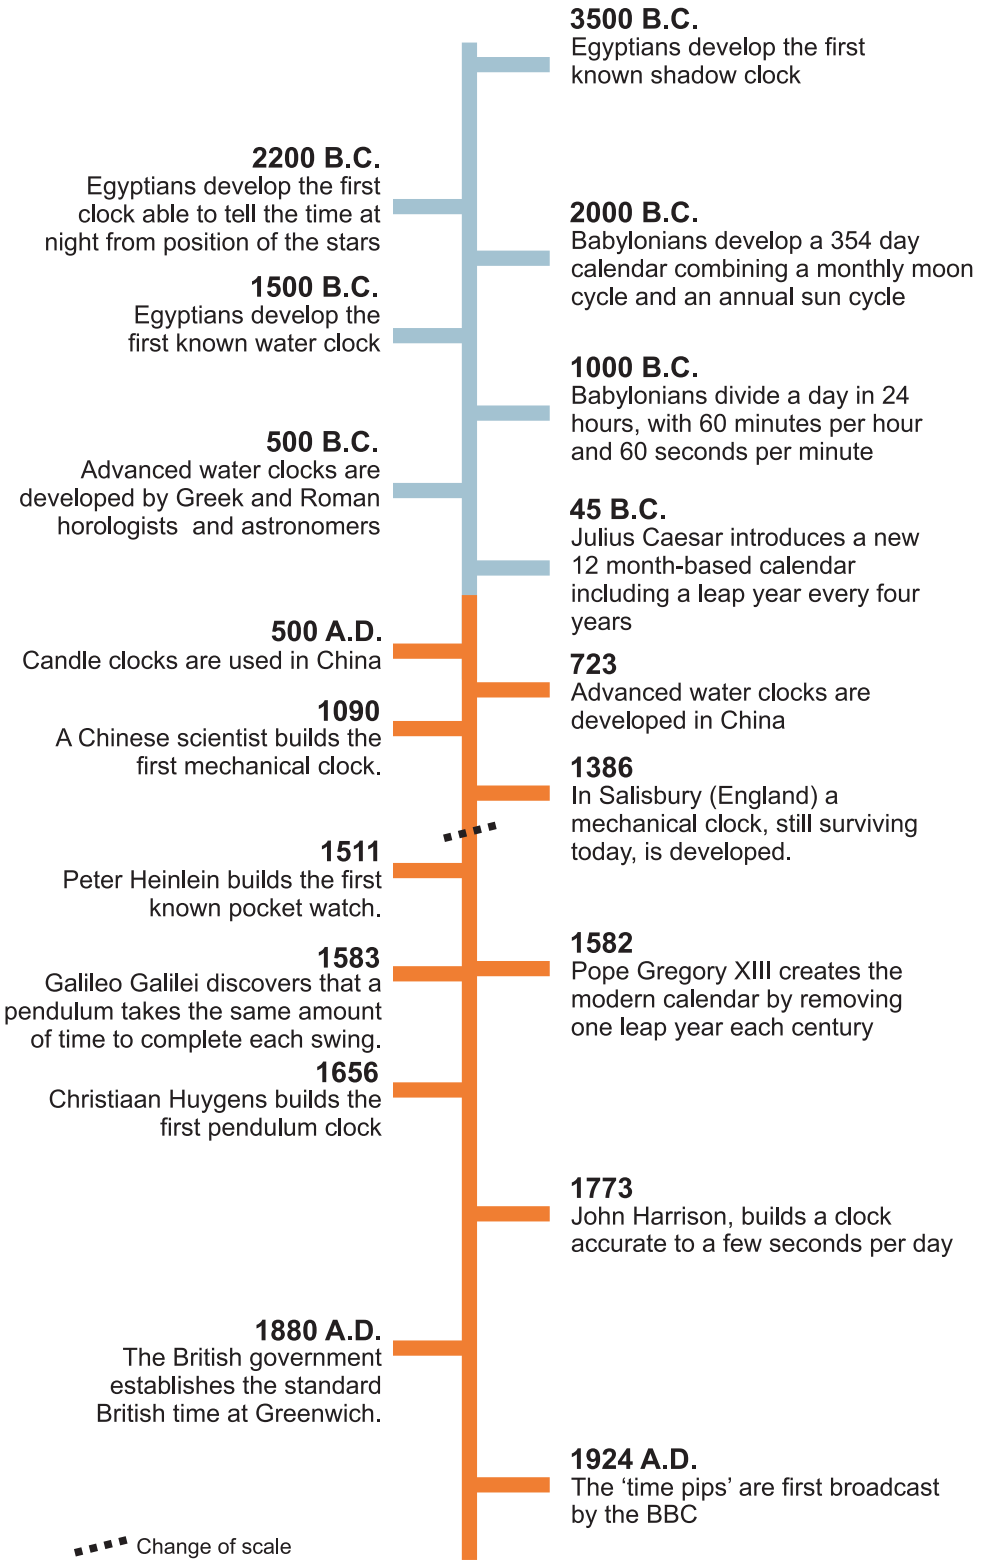

Figure 1.1: The defining moments of the art of timekeeping from 3500 B.C. to 1924.

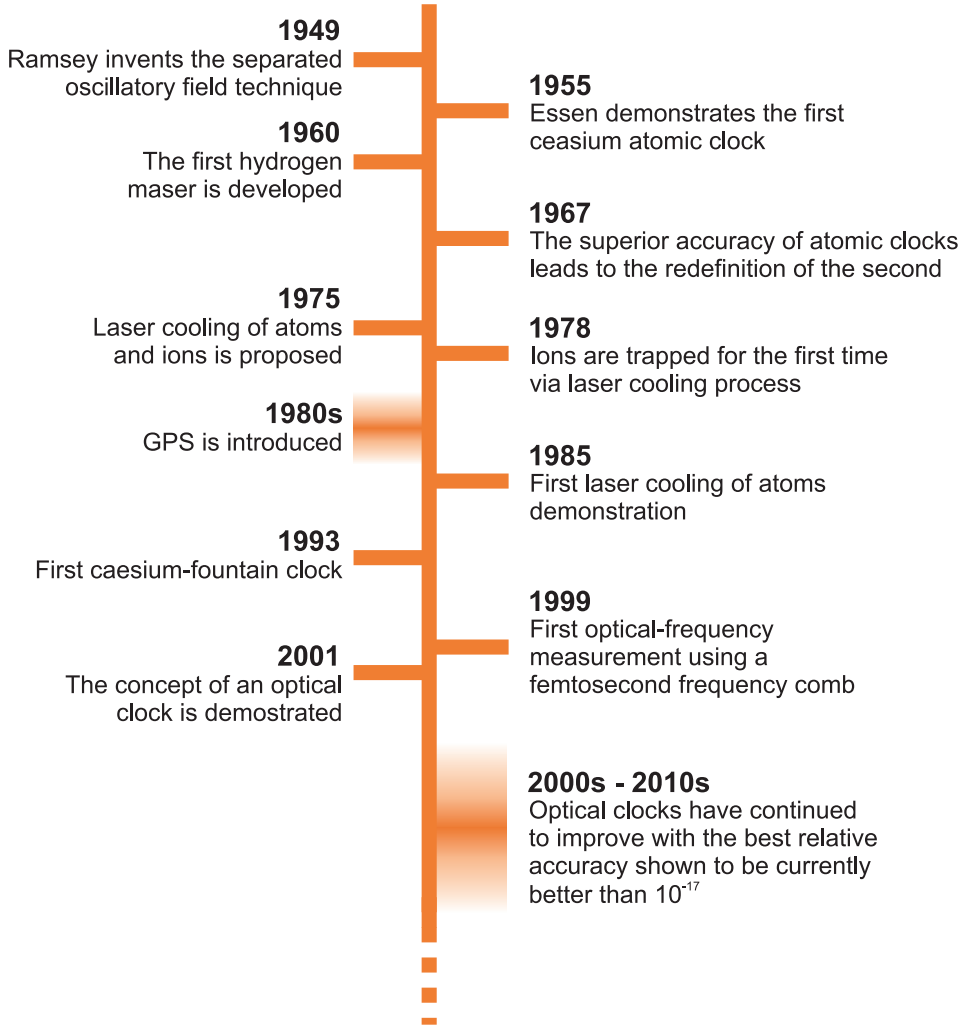

Figure 1.2: The defining moments of the modern era of timekeeping starting with Ramsey in 1949.

In Figure 1.3 two families of atomic clocks are shown. Both use discrete atomic energy levels as a reference, but the frequency of the atomic transition used resides in two different parts of the electromagnetic spectrum. Essen's clock was based on a microwave transition and further developments of this type of clock have allowed its accuracy to be improved by over 5 orders of magnitude, with the current generation of caesium atomic fountain clocks capable of an accuracy of a few parts in 10^{16} over 1 day (Jefferts et al., 2007; Li et al., 2011; Guéna et al., 2012; Weyers et al., 2012). When lasers became widely available, scientists started to investigate whether it was possible to use optical rather than microwave transitions. As optical frequencies are four to five orders of magnitude higher that microwave frequencies, it should be possible, at least in principle, to subdivide time into thinner "slices" and therefore to measure it more accurately. The first measureament of an optical transition referenced to the caesium standard, using a phase-coherent microwave-to-optical chain, was performed in 1996 at the Physikalisch-Technische Bundesanstalt (PTB) (Schnatz et al., 1996). The accuracy

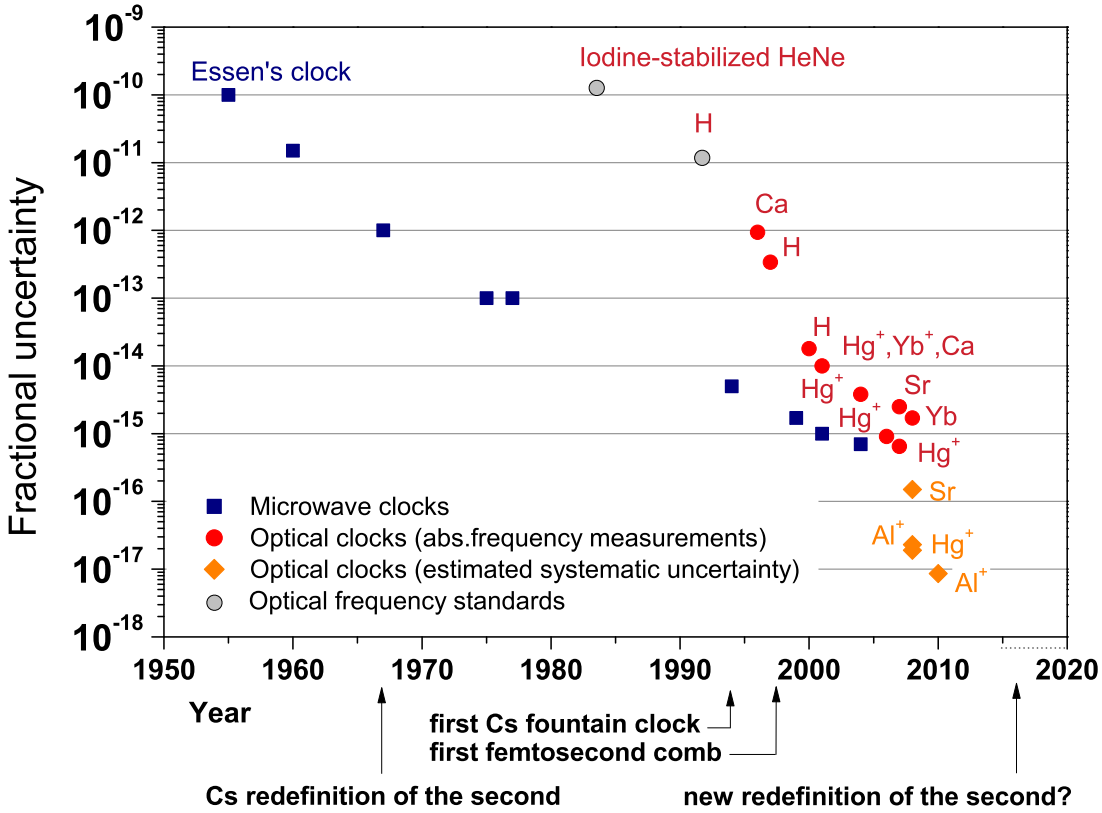

Figure 1.3: Improvements in the accuracy of microwave and optical clocks in the last 50 years. The introduction of the femtosecond comb technique has sped up enormously the development of optical clocks in the last 10 years. A new redefinition of the second might take place in the next decade or so. (Modified figure from the original by Helen Margolis (NPL)

of this optical clock was orders of magnitude worse than the best microwave atomic clocks available at the time, but the pioneering experiment paved the way for further developments of clocks of this type. The first optical clock emplyoing a femtosecond optical frequency comb was developed in 2001 by NIST (Diddams et al., 2007). If the improvement of microwave atomic clocks has been impressive, that of optical clocks is even more remarkable. In less than 15 years their accuracy has increased by almost 5 orders of magnitude (Gill, 2011).

1.4 How accurate is a clock? The need for comparison

At the beginning of the 20th century, when Londoners wanted to check whether their pocket watch was correct, they would compare it with Big Ben's clock. They could do so as they trusted the tower's clock to be more accurate than their pocket watch. Every time we wish to establish the accuracy of a clock, we need to compare it to another of better accuracy. But how do we assess the accuracy of the best clock on Earth? As

we do not have anything better to compare with, the solution is to build two identical clocks and see by how much they differ as time passes. The measured accuracy will be, in the best case, that of the less accurate of the two clocks. This is the reason why in all laboratories where clocks are developed, a minimun of two are always built. However, the ultimate test to assess their accuracy is to compare them with others developed independently in other laboratories. Clock comparisons within the same laboratory are certainly required for initial assessments, but they might be subject to the same environmental effects and human errors. Inter-laboratory comparisons provide the best possible confidence test of the results.

1.4.1 Comparison of optical clocks in Europe

Optical clocks are mainly being developed in three locations around the world: Europe, the USA and Japan. I will consider here only Europe as it has the highest concentration of optical clocks, making it the most probable location for most of the future remote optical clock comparisons. In Figure 1.4, the location of the European laboratories where optical clocks are being developed are shown. The average distance between laboratories is many hundreds of kilometres. Currently, remote clock comparison is achieved by means of microwave satellite links. However, the stability of the link is at best at the 10^{15} level over one day (Bauch et al., 2006) and it is therefore insufficient to preserve the stability and accuracy of the optical clocks. The technique is mainly limited by atmospheric effects, which are extremely difficult to compensate at the level of accuracy required. Some improvements in the long term stability can be expected in the next few years, while the short term stability is likely to remain insufficient for most state-of-the-art metrology applications.

With the above considerations in mind, in the last few years some research groups have started investigating the possibility of performing improved remote clock comparisons by propagation of laser light, either continuous wave or pulsed, over fibre networks. The results achieved so far already show this solution as a viable alternative to satellite links (Jiang et al., 2008; Williams et al., 2008; Musha et al., 2008; Grosche et al., 2009; Lopez et al., 2010; Marra et al., 2011; Predehl et al., 2012). At the time of writing, a large European Metrology Research Programme[4] (EMRP)-funded project between many European measurement insitutes is planned to start in the summer of 2012. The aim is to develop the transfer techniques and the practical solutions required to implement time and frequency transfer between laboratories which are separated by distances as long as 1500 km. The development of a pan-European optical network will undoubtedly further speed up the development of the future generation of optical clocks as well as enable new experiments in fundamental physics and other scientific fields to be carried out. One of the tasks NPL will undertake in the European project is to investigate the

[4]The European Metrology Research Programme is a metrology-focused programme that coordinates and facilitates national research programmes to accelerate innovation and competitiveness in Europe.

Figure 1.4: Geographical distribution of the European laboratories currently developing optical clocks.

transfer of time using frequency combs, which will build on the work presented in this thesis.

1.5 Content of this thesis

This thesis is organized into 6 chapters as follows:

Chapter 2 - Background This chapter gives an overview of the basics of frequency metrology, from the concept of accuracy and stability to the measurement techniques. It also discusses the fundamentals of optical frequency combs and frequency stabilization techniques.

Chapter 3 Techniques for frequency transfer over optical fibe - In this chapter an overview of the operating principle and test results of the frequency transfer techniques investigated to date is presented.

Chapter 4 - Optical to microwave conversion In this chapter the noise associated with the optical-to-microwave conversion of the pulse repetition rate of a mode-locked laser, a crucial aspect in all the experiments described in this thesis, is discussed and measured experimentally. An experiment showing how the detected microwave signal changes when the pulse width is changed is also presented.

Chapter 5 - Comb-based microwave transfer over 50 km of spooled fibre In this chapter the transfer of a 1.5 GHz microwave frequency by propagating an optical frequency comb over 50 km of spooled fibre is presented. Unlike to previous experiments performed elsewhere, this experiment measures the actual stability delivered at the "user"' end, whilst also demonstrating improved transfer stability.

Chapter 6 - Microwave transfer over an installed fibre In this chapter the transfer of a microwave frequency using an optical frequency comb is achieved over an 86 km-long link of the JANET-Aurora dark fibre network. The experiment demonstrated transfer stabilities comparable to those measured on spools in the lab, whilst also showing for the first time the long term performance. Differently from the previous experiment, the pulse train was propagated in both directions within the same fibre and optical amplifiers were employed at each end.

Chapter 7 - Optical phase detection and test of the full comb structure In this chapter an optical, rather than a microwave, phase detection technique is employed to cancel the phase noise introduced by the fibre. For the first time, the stability and accuracy of the individual optical modes, rather than their spacing (repetition rate), is measured. The experiment shows that is it possible to transfer many thousands of optical modes at a level comparable to that achievable by single carrier transfer techniques.

Chapter 8 - Conclusions In this chapter conclusion are drawn and future work and possible applications of the investigated techniques are discussed.

Each of the experiments described in chapters 3,4 and 5 have resulted in a peer-reviewed publication, as outlined at the beginning of each chapter.

Chapter 2

Background

In the previous chapter I have repeatedly used the terms accuracy and stability. Although their physical meaning differs substantially, these terms are often misused and it is worth briefly recalling these concepts here. I will also discuss the concepts of phase noise and frequency stability and the techniques that allow their measurement, as they will be required to understand the experimental work performed.

2.1 Accuracy and stability

The *accuracy* is the degree of conformity of a measured or calculated value to its *true value*. This definition might seem recursive as, in the case of a physical quantity, the true value is itself the result of a measurement. In practice, the true value of a physical quantity is generally agreed internationally as a result of multiple independent measurements performed by different laboratories around the world.

The *stability* describes the variation of a physical quantity over time. The difference between accuracy and stability can be better illustrated with the example shown in Figure 2.1, where the frequency of an oscillator is analysed. As shown, it is possible for an oscillator to be stable but not accurate and *vice versa*.

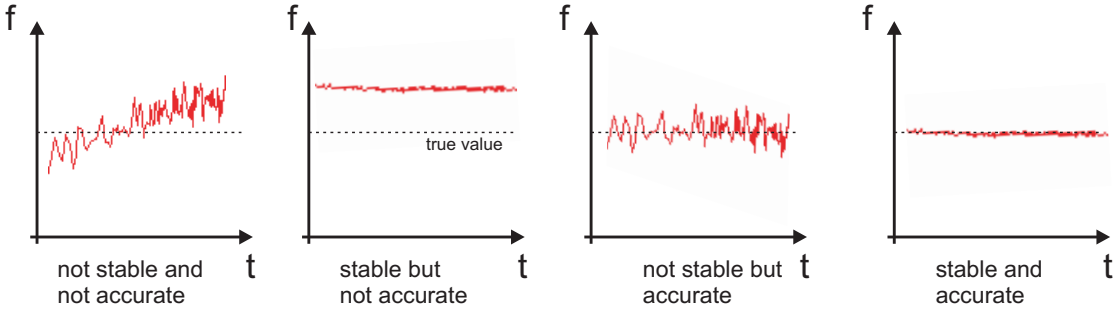

Figure 2.1: Illustration of the difference between accuracy and stability

2.2 Characterization of the frequency fluctuations

The signal produced by an ideal oscillator is an ideal sinuoid that can be written as

$$x(t) = A_0 \cos(\omega_0 t + \phi) \tag{2.1}$$

where A_0 is the peak amplitude of the sinusoid, ω_0 is its angular frequency and ϕ is a constant phase shift with respect to an arbitrary zero phase point. This equation applies to any type of oscillator. In a pendulum A_0 will be the maximum swing of the weight from the resting position whilst in an electronic oscillator it will be the maximum output voltage. In the real world, a number of physical processes can affect the amplitude and phase of the pure sinusoid, making these quantities randomly fluctuate with time. These physical processes are usually of enviromental or electronic nature in modern non-mechanical oscillators such as those that will be discussed in this thesis. These perturbations are generically called "noise", but we will see later how many types of noise can be present in nature and how much information can be extracted from the measurement of their statistical properties. The "noisy" signal of a real oscillator can be described as

$$x(t) = A_0[1 + \alpha(t)] \cos(\omega_0 t + \xi(t) + \phi) \tag{2.2}$$

where $\alpha(t)$ and $\xi(t)$ are random variables describing respectively amplitude and phase fluctuations. The concepts of amplitude and phase fluctuations can be particularly well illustrated with the *phasor* representation depicted in Figure 2.2. In order to characterize

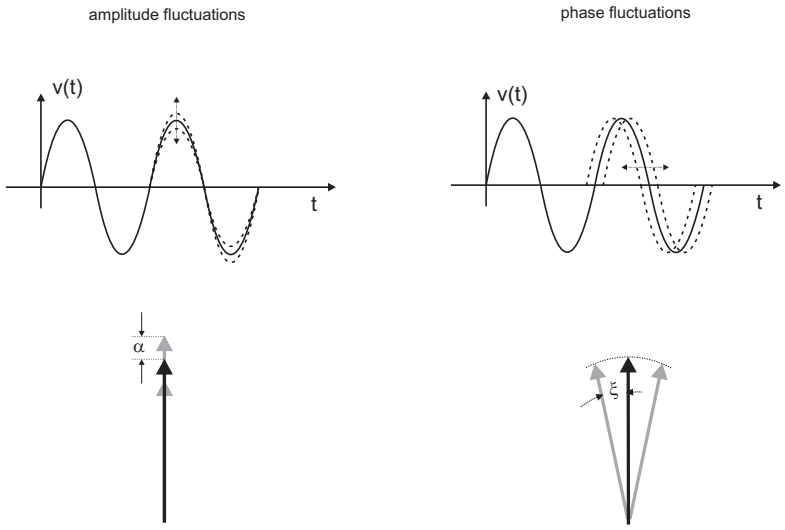

Figure 2.2: Phasor representation of amplitude and phase noise

the "quality" of an oscillator, the fluctuations of the amplitude and phase over time have to be measured and some form of averaging needs to be computed. This characterization can be performed either in the time domain or in the frequency domain. For phase fluctuations measurements, one usually refers to it as *frequency stability* for the time domain characterization, and *phase noise* when the characterization is performed in the frequency domain. It should be noted however that both frequency stability and phase noise measurements are two ways to describe the very same statistical properties of an oscillator.

2.2.1 Phase noise and frequency stability

The reason why two methods are used to describe the same properties is mainly due to practical constraints. The resolution of a frequency counter is directly proportional to the measurement (sample) time and is generally insufficient to measure the very good short term stability achieved by modern oscillators. The analysis of the oscillator's phase in the frequency domain is therefore used in this case. Generally, if the stability of an oscillator over time intervals above 1 s has to be characterized, frequency stability measurements are used. If fast fluctuations of the frequency are of interest, phase noise measurements are instead preferred.

Frequency stability A frequency stability measurement provides a statistical description of the variation in the number of occurrences of a specific event over a certain time. For example, a possible frequency stability measurement is to count the number of zero crossings of a sinusoidal signal taking place over one second, to repeat the measurement every second and then to calculate some kind of average of all the measurements. For a pure sinusoid of frequency of 1 MHz, *exactly* 10^6 zero crossings will always occur in 1 s. The variation from one measurement to the next, and to any other, will be zero. The frequency stability in this case would be *infinite* (and conversely the instability would be zero). If the phase of the sinusoid fluctuates with time, it is possible that the number of zero crossings will be lower or higher than 10^6 in 1 s from time to time. The variation from one measurement to the next, and to any other, will not be zero in this case. The frequency measurement will provide the *average value* of the frequency of the oscillator over a measurement time of 1 s.

The frequency stability of an oscillator is often expressed in fractional terms by dividing the absolute frequency variations by the oscillator frequency. This normalization makes it easy to compare the stability of oscillators working at different frequencies. For example, both cryogenic sapphire oscillators and ultra-low expansion (ULE)-cavity stabilized lasers (Young et al., 1999; Stoehr et al., 2006; Notcutt et al., 2006; Ludlow et al., 2007; Webster et al., 2008; Dubè et al., 2009; Lodewyck et al., 2010; Millo et al., 2009; Webster and Gill, 2011; Jiang et al., 2011; Watabe et al., 2006; Hartnett et al., 2012) deliver a

stability at the 10^{-15} level of better, but their frequency differs by at least a factor of 10^4. The frequency stability can be measured over different time intervals generating a stability-vs-time plot like that shown in Figure 2.3. The plot tells us that that the stability of the frequency of the oscillator is better over intervals of 100s s than over intervals of 1 s and that beyond 100 s it does not improve any further. We will see later how from the steepness of the slopes it is possible to make assumptions on the dominant noise process occurring at each given time scale.

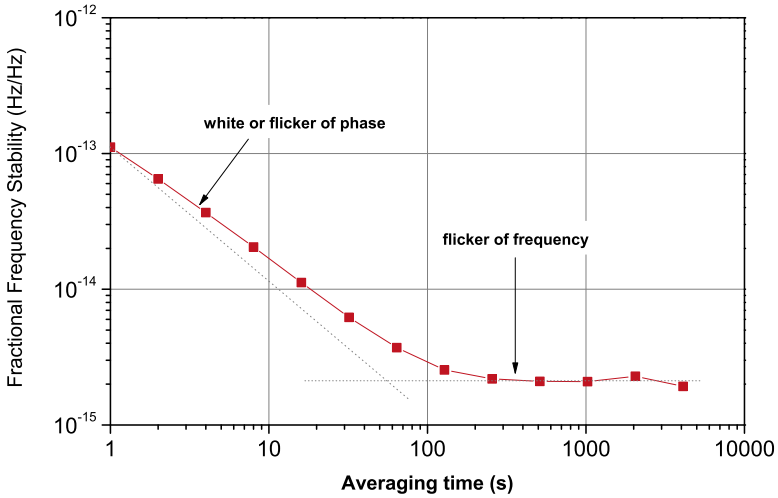

Figure 2.3: A typical frequency stability plot. From the slopes it is possible to extract useful information about the dominant noise process at each time scale.

Phase noise The outcome of a phase noise measurement is the power spectral density (PSD) of the phase fluctuations $\phi(t)$ and it is defined as

$$S_\phi(f) = |\Phi(jf)|^2 \tag{2.3}$$

where $\Phi(jf)$ represents the Fourier transform of $\phi(t)$. It is usually calculated using a Fast Fourier Transform (FFT) analyser or a spectrum analyser from the phase samples. The PSD can be interpreted as a statistical description of how often the oscillator frequency momentarily assumes values other then its nominal frequency. In the simple case of flicker phase noise ($1/f$), the frequency assumes all values with a probability inversely proportional to the offset frequency[1]. Figure 2.4 shown a typical phase noise plot.

The phase noise is generally represented as $S_\phi(f)$, although the single sideband phase noise $\mathcal{L}(f)$ is also often used:

$$\mathcal{L}(f) = \frac{1}{2} S_\phi(f). \tag{2.4}$$

[1]In a real oscillator, at small offset frequencies the carrier phase noise increases as $1/f^3$, whilst the flicker phase noise is generally due to buffer amplifiers used to increase the output power.

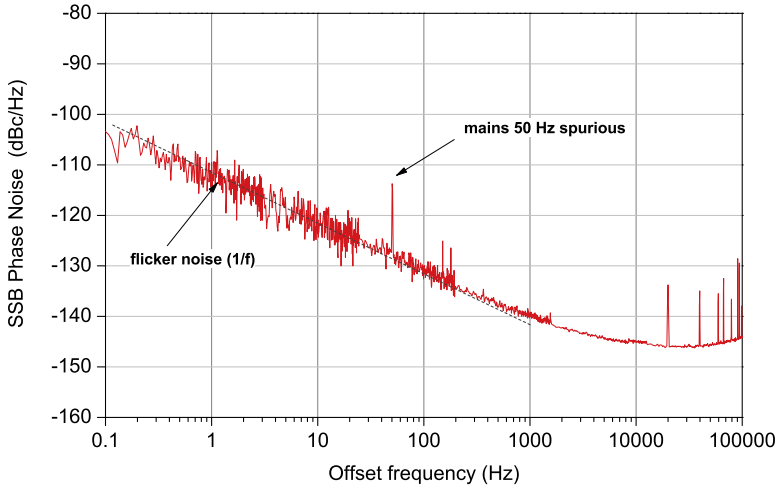

Figure 2.4: A typical phase noise plot showing flicker noise ($1/f$).

$S_\phi(f)$ is expressed in dBrad^2/Hz, while $\mathcal{L}(f)$ is expressed in dBc/Hz which has the meaning of "dB with respect to the carrier in a 1 Hz bandwidth".

2.2.2 Allan variance

In the frequency measurement example described earlier, the frequency samples were collected and then averaged. The most common time-average used to characterize the frequency stability of oscillators is a two-sample variance called the Allan variance (D. B. Sullivan and Walls, 1990; Sullivan et al., 1990; Rutman, 1978; Vig, 1999), defined as

$$\sigma_y^2(\tau) = \frac{1}{2(M-1)} \sum_{k=1}^{M-1} [\bar{y}_{k+1} - \bar{y}_k]^2 \tag{2.5}$$

where M is the total number of samples and $\bar{y}_{k+1} - \bar{y}_k$ is the difference between two consecutive samples of the fractional frequency average over the observation time τ. The Allan variance is a powerful tool that allows the noise processes affecting the system under test to be distinguished. For each noise process the Allan variance produces a different stability slope (in a log-log plot) with respect to the integration time. The only two noise processes that cannot be distinguished are the white and flicker of phase. A modified Allan variance was developed in order to separate these noise processes, although the standard Allan variance remains the most commonly used. Most of the time, the square root of the Allan variance $\sigma_y^2(\tau)$, which takes the name Allan *deviation*, is used as it has more physical meaning being expressed in the same units as the measured data rather then the units squared. The different Allan deviation slopes arising from

the five most commonly found noise processes, white and flicker of phase, white, flicker[2] and random walk of frequency are shown in Figure 2.11.

2.2.3 Conversion between phase noise and frequency stability

The spectra of most noise processes can be modelled by a power-law function, as depicted in the top half of Figure 2.5. From the phase noise spectrum it is possible to calculate[3] the corresponding Allan deviation using the formule found, for example, in (Sullivan et al., 1990; Dawkins et al., 2007). The calculation of the phase noise spectra from Allan deviation is not always possible, as both white and flicker phase noise have the same dependence on τ in the Allan deviation. It is worth highlighting that the x-axis in the Allan deviation plot is time while in the phase noise plot it is frequency. As we will see later, modern frequency counters generate other types of variance rather than the Allan variance. Nonetheless, the power-law model is still valid and the conversion from phase noise is possible. The noise processes most often encountered in the experiments described in this thesis are white and flicker of phase and flicker of frequency.

2.2.4 Phase noise measurement techniques

Among the several phase noise measurement techniques possible I will restrict my description to the case where two signals of identical frequency are readily available for phase comparison. This will in fact be the case for the experiments described in this thesis. The relative phase between two signals of the same frequency can be measured using an RF or microwave mixer. A mixer multiplies the input signals and, if the mean phase difference between two sinusoidal inputs is maintained at $\frac{\pi}{2}$ (under this condition the signals are said to be in *phase quadrature*), then the output voltage of the mixer is

$$V_{\text{out}}(t) \propto \cos\big(\omega_0 t + \xi(t)\big) \cos\big(\omega_0 t - \frac{\pi}{2}\big) = \frac{1}{2}\sin\big(2\omega_0 t + \xi(t)\big) - \frac{1}{2}\sin\big(\xi(t)\big) \tag{2.6}$$

For small phase changes, $\sin(\xi(t)) \sim \xi(t)$. The component at $2\omega_0$ can be easily removed with a low pass filter. We can rewrite the relationship between the output voltage of the mixer and the phase difference between the input signals by defining the mixer phase sensitivity k_ϕ

$$k_\phi = \frac{V_o}{\xi} \tag{2.7}$$

which has units V/rad. The value of k_ϕ is dependent on the power applied to the LO and RF ports of the mixer. The calibration of k_ϕ is thus required each time a measurement

[2]The flicker of frequency is equivalent to random walk of phase.

[3]A very useful Excel spreadsheet for the conversion of phase noise into frequency stability can be found at www.wenzel.com. Wenzel Assiciates Inc. is an American company producing some of the lowest phase noise oscillators in the 50-100 MHz range.

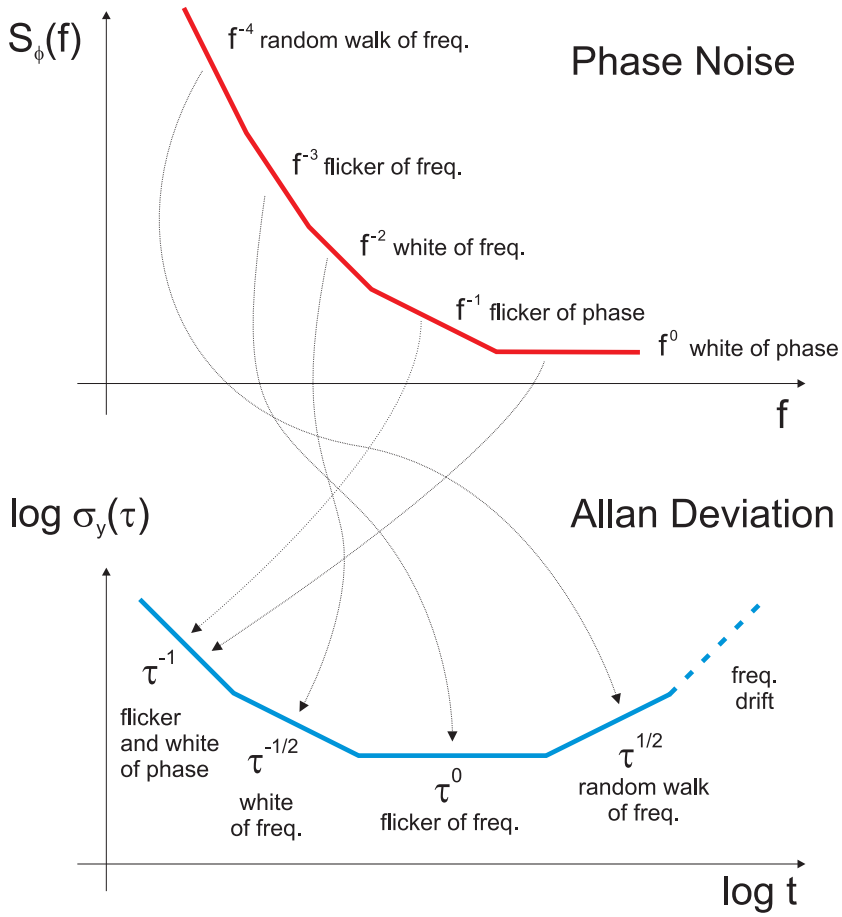

Figure 2.5: Relationship between phase noise power-law spectrum and Allan deviation.

is taken if the input power is changed. In order to obtain the power spectral density of the phase fluctuations, the output voltage of the mixer is amplified and sent to a Fast Fourier Transform (FFT) analyser.

Despite the apparent simplicity of this phase noise measurement method, it turns out not to be trivial in practice. Many environmental factors and subtle phenomena can corrupt the measurement in non-obvious ways. Phase noise measurement of extremely high stability signals is particularly difficult as power fluctuations of the input signals, power supply noise and thermal fluctuations can overwhelm the actual phase noise to be measured.

2.2.5 Practical aspects of measuring frequency stability and phase noise

In this short section I wish to discuss some practical aspects of frequency stability and phase noise measurement of state-of-the-art frequency references. With this insight I

hope to provide a better feeling for the practical problems that the scientist faces in the lab and that are hard to envisage from the theoretical background.

Frequency stability measurements A direct frequency stability measurement can usually be achieved by counting the zero crossings of the signal under test with a frequency counter. This is straightforward when the signal under test has a frequency lower than a few hundred MHz and the fractional instability to be measured is expected to be higher than 10^{-11} or 10^{-12}, which is the resolution of modern frequency counters. For state-of-the-art frequency metrology this level of instability is orders of magnitude worse than what is desired and methods for enhancing the measurement resolution need to be implemented. A selection of these techniques are discussed below. Some modern frequency counters implement advanced algorithms for resolution enhancement during the frequency counting process[4]. Whilst these type of counters have been very widely used for many years, it was found only recently that applying the Allan deviation formula to the data sets from these counters does not give the true Allan deviation but another type of two-sample variance which can differ substantially for certain noise processes. This undesired effect was highlighted by a remarkable reverse-engineering work by Rubiola (Rubiola, 2005) and then perfected by Dawkins (Dawkins et al., 2007) in a later paper. These counters are usually referred to as Λ-type and the two-sample variance calculated from their data sets has been named "Triangle" variance. In opposition, non-resolution enhanced frequency counters are usually referred to as Π-type. The Greek letters have been chosen to graphically represent the weighting function applied to the frequency samples and more about this can be found in the references provided.

Another problem afflicting most frequency counters is the dead-time between frequency readings. The dead-time affects the calculation of the Allan variance, showing up as a change in slope for some noise processes (Barnes et al., 1990). Only a small selection of recent counters can operate in zero-dead-time mode[5].

Spurious signals and mains leakage can also affect the frequency counting process usually appearing as bumps in Allan deviation plots. This problem becomes particularly tedious when short measurement intervals τ are used as they become comparable with the period of the spurious signals (for example, τ is 20 ms for the 50 Hz mains frequency).

Lastly, the resolution of the counter needs to be tested as it is a function of a number of parameters such as the frequency and slew rate of the input signal and the timing reference used. In some cases the resolution degradation approach an order of magnitude. If the expected stability of the signal under test is only two or three time worse than

[4]Counters of this types are the popular Agilent 53132A and the Pendulum CNT-90/91. The formerly very popular Stanford Research SR620 is instead a Π-type counter. It is useful to know that the Agilent 53132A counter can work as a Π-type counter when it is externally armed.

[5]Examples of frequency counters operating in zero-dead-time mode are the Pendulum CNT-91 and those from the small German company K+K.

the counter's resolution (and this can often be the case for state-of-the-art frequency metrology), tests will be required to make sure of the validity of what is being measured.

The combination of the problems described above often make the interpretation of the results more laborious than one might initially expect. Whenever possible, a frequency stability measurement should be complemented by a phase noise measurement and the consistency of the results should be cross-checked.

Phase noise measurements The measurement of phase noise is generally even more prone to errors than a frequency stability measurement performed with frequency counters. Phase noise measurements are in fact most often performed with power-sensitive devices such as analog phase detectors and mixers. The power sensitivity forces a device-specific calibration, rendering the measurement more laborious and prone to human errors and approximations. Digital phase detectors are much less sensitive to power fluctuations as they are based on digital devices triggered by zero-crossing of the input signal, where a change in power leads to a much smaller change in the detected phase than in the analog case. However, they are generally too noisy for most RF and microwave applications. In state-of-the-art measurements, even the noise of analog phase detectors or mixers can also be a concern.

Techniques for the enhancement of the measurement resolution The effect that some basic operations such as a division or a sum have on the measurement of frequency stability and phase noise is described here. Understanding these basic concepts enables us to use them for resolution enhancement techniques.

1. **Algebrical sum (down- or up-conversion)** A frequency mixer generates the sum and the difference of the input frequencies. The process is called "downconversion" when two frequencies are mixed together to produce a lower frequency (beat note) and "up-conversion" in the opposite case. The absolute frequency deviation of either of the two inputs will be transferred directly onto the beat note. For example, let us assume a signal of frequency 100 MHz which fluctuates on average by 1 Hz (hence its fractional stability is 10^{-8}) is beat against a substanstially more stable frequency of 101 MHz. The resulting beat will be at 1 MHz and will fluctuate by the same absolute amount as the 100 MHz signal. However, given that the beat note is lower in frequency, the fractional frequency stability of this signal will be 10^{-6}. We can use this as a leverage effect in order to measure high frequency signals with enhanced resolution. If a frequency counter is able to measure a signal of 1 MHz with a resolution of 10 digits (10^{-10}) and we need to measure the stability of a 10 GHz signal, a resolution of 10^{-14} can be achieved by downconversion, provided that a more stable source can be used to produce the beat. The same effect of course can be used to measure phase noise. The

process of downconversion also happens in a photodiode when two or more optical frequencies are present. The resulting beat note is an RF or microwave frequency whose stability can easily be counted by a common frequency counter. Given that the ratio between the optical frequencies and the rf frequency is usually higher than 10^5, resolutions exceeding 10^{-15} can easily be achieved.

2. **Division and multiplication** The process of dividing or multipling a frequency has no effect on the fractional frequency stability of the generated signal. The phase fluctuations are instead accordingly divided or multiplied. The division can be used in some cases to reduce the phase fluctuations of a noisy signal down to the linear range of a phase detector. A typical example of when a divider can be used is the phase noise measurement of optical beats. The phase noise of the resulting signal is scaled as N^2 where N is the division or multiplication ratio. For example, a signal frequency divided by 100 will have a phase noise 40 dB lower than that at the input of the divider.

The properties of the operations described are summarized Table 2.2.5.

Operation	output fractional stab.	output phase noise
Up-conversion	higher by conversion factor	same as input
Down-conversion	lower by conversion factor	same as input
Division by N	same as input	lower by N^2
Multiplication by N	same as input	higher by N^2

2.2.6 Timing jitter

In some cases it is useful to express the phase fluctuations as an average timing jitter, specified over a certain time scale and noise bandwidth, rather than a power spectral density. The timing jitter of a sinewave can be visualized as the fluctuations of the zero-crossings (I tend to think of it as a measure of how "blurred" the zero-crossings are). Typical applications where this quantity is preferred are generally those involving event syncronization, such as optical communications, radar and radioastronomy. The root mean square timing jitter, expressed in units of seconds, is calculated as

$$T_{\mathrm{rms}} = \frac{1}{2\pi f_o}\sqrt{\int_{f_l}^{f_h} S_\phi(f)\,\mathrm{d}f}, \tag{2.2.8}$$

where f_l and f_h are the bandwidth's lower and upper limits respectively, S_ϕ is the double-sided phase noise and f_o is the frequency of the signal under test. It is instructive to calculate the timing jitter expected from electronic and optical devices used for frequency metrology:

Device	Freq. (Hz)	Jitter (fs)
Ultra stable crystal oscillator	5×10^6	320
Room temperature sapphire oscillator	10×10^9	45
Cryogenic sapphire oscillator	9×10^9	10-20
ULE-stabilized laser	4×10^{14}	0.1-0.4

all value are calculated over a bandwidth extending from 1 Hz to 100 kHz.

2.3 Fundamentals of optical frequency combs

An optical frequency comb is a series of equally spaced discrete frequencies extending over a certain optical bandwidth. In an ideal comb, all the modes are in phase. This means that, despite all modes having different frequencies, there will be a moment in time where the electric field of each optical mode add constructively and a high intensity pulse is generated. This concept can be better understood by observing how the superposition of an increasing number of sinusoids with harmonic frequency progression but the same initial phase give rise to an increasingly pulse-like intensity profile. In real optical frequency combs, the perfect *in-phase* relationships between all the modes is corrupted by the chromatic dispersion of the optical materials used in the pulse-generation. The dispersion causes the group velocity to differ from the phase velocity of the optical carrier. This imperfection manifests itself in the frequency domain as a frequency offset superimposed on all the optical mode frequencies as depicted in Figure 2.7. This offset is usually called the carrier-envelope offset frequency f_{CEO}. The frequency of the optical carrier of the pulse is the average[6] frequency of all the optical modes.

Two parameters are therefore needed to fully characterize each mode of an optical frequency comb, the carrier-envelope frequency f_{CEO} and the frequency spacing between the optical modes f_r (repetition rate):

$$f_n = n f_{\mathrm{r}} + f_{CEO} \tag{2.3.1}$$

where n is the mode number and it is usually of the order of 10^5 or 10^6.

2.3.1 Optical frequency comb stabilization

The carrier-envelope offset frequency and the repetition rate are usually too unstable for state-of-the-art frequency metrology and they need to be stabilized. A description of the most common comb stabilization techniques follows here.

[6] A weighted average might be required if the comb intensity of the modes or the dispersion is not symmetric with respect to the center frequency (or both).

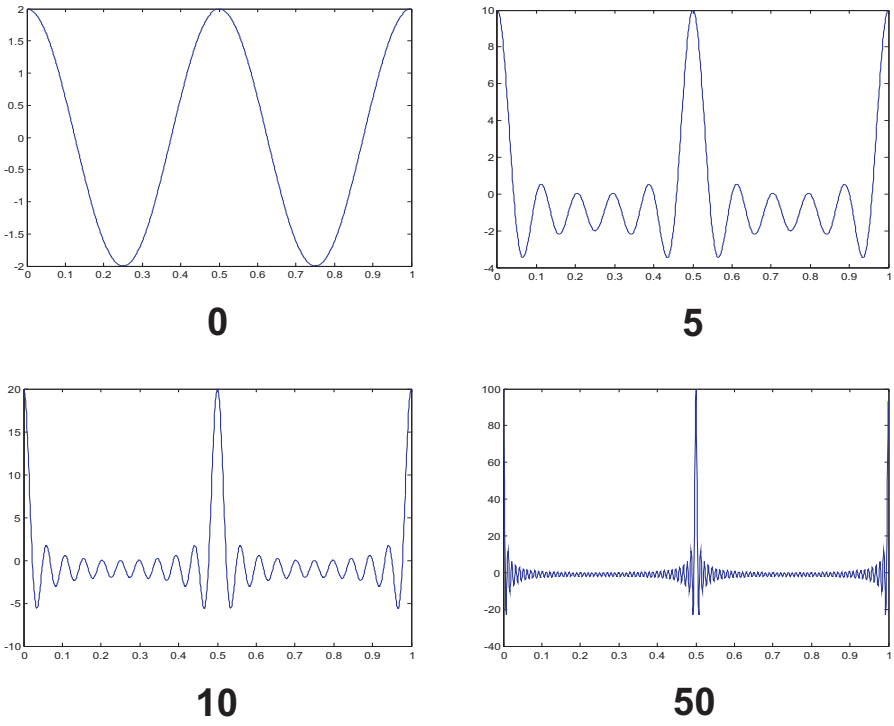

Figure 2.6: Generation of a pulse from in-phase superposition of N sinusoids.

Carrier-envelope offset frequency stabilization The carrier-envelope offset frequency can be measured using the technique described in (Jones et al., 2000; Holzwarth et al., 2000) and shown in Figure 2.8 which is usually referred to as f-$2f$ interferometry. A heterodyne beat is constructed between the high frequency portion of the optical spectrum and the frequency-doubled low frequency portion. Each comb mode of the frequency-doubled spectrum can be described as

$$2f_n = 2nf_{rep} + 2f_{CEO} \tag{2.3.2}$$

whilst those of the high frequency portion are described by:

$$f_{2n} = 2nf_{rep} + f_{CEO} \tag{2.3.3}$$

by extracting the beat between these two frequency components with a photodiode allows f_{CEO} to be measured:

$$2f_n - f_2 n = 2nf_{rep} + 2f_{CEO} - 2nf_{rep} + f_{CEO} = f_{CEO} \tag{2.3.4}$$

In an actual implementation, many pairs of optical modes are used to maximize the signal-to-noise ratio of the to extracted f_{CEO} as the power available in each mode is usually small. The f-$2f$ interferometric method requires that the optical comb extends over at least an octave. Full octave frequency combs directly generated by femtosecond

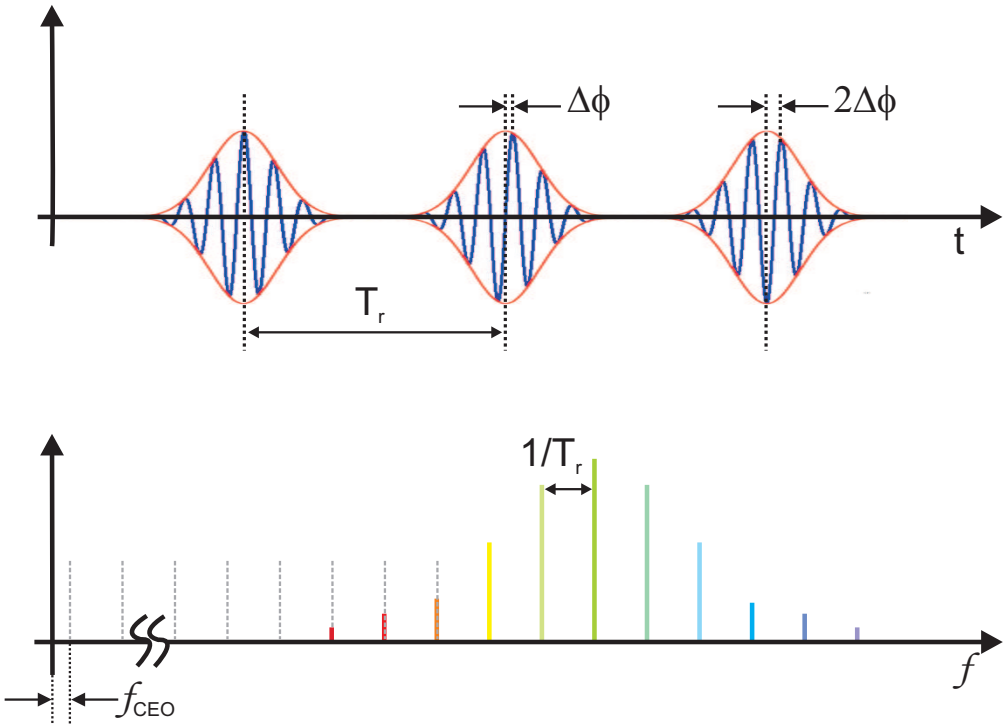

Figure 2.7: Illustration of the origin of the carrier-envelope frequency offset found in real optical frequency combs. The offset originates from the difference between phase and group velocity within the laser cavity caused by chromatic dispersion.

lasers have been developed (Ell et al., 2001; Fortier et al., 2003) but most often the octave spectrum is obtained from a relatively narrow bandwidth comb (typically a few tens of nm full width half maximum) by supercontinuum generation using a microstructured fibre. The frequency comb used for all the experiments described in this thesis does indeed employ a microstructured fibre to achieve the spectral width required for the f-$2f$ detection. Other f_{CEO} detection schemes exist that require only 2/3 of an octave or less, but they are more complex to implementat (Morgner et al., 2001; Ramond et al., 2002).

f_{CEO} is usually an RF or microwave frequency which can be phase compared to a stable oscillator. The error signal arising from the comparison can be then used to control, for example, the pump power intensity of the laser which in turn changes f_{CEO} (Reichert et al., 1999; Holman et al., 2003). It is worth noticing that the fractional stability of the reference RF or microwave oscillator does not need to be particularly high. If, for example, a stability of 1 Hz is wanted on the optical modes of the frequency comb (corresponding to 5×10^{-15} for a frequency comb of center frequency 200 THz), the reference oscillator will need to be stable only at the level of 10^{-7}, if the f_{CEO} is 10 MHz, and 10^{-9} if f_{CEO} is higher than a few 100 MHz. Most commercial synthesizers are able to deliver such stablity.

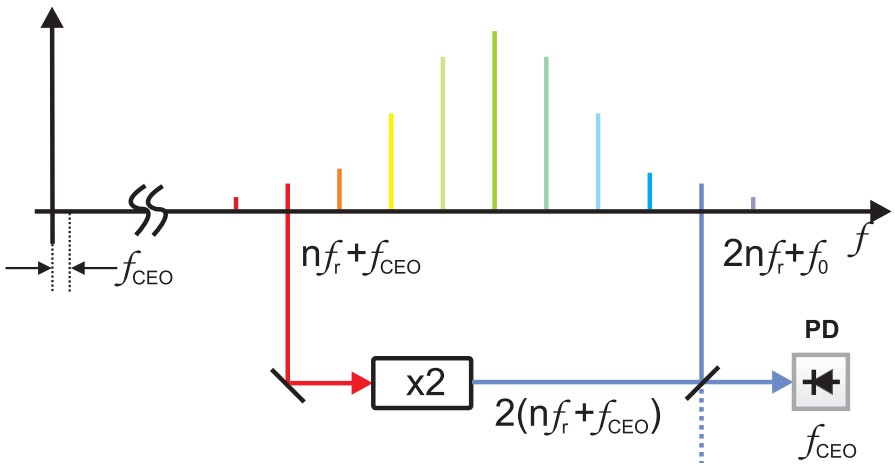

Figure 2.8: Diagram illuastrating the principle of the f-2f interferometer that allows the carrier-envelope offset frequency to be measured in an optical frequency comb.

Repetition rate stabilization The detection of the repetition rate is much more straightforward than that of f_{CEO} and can be accomplished using a simple photodetector. The repetition rate depends primarily on the cavity length of the laser resonator and it can usually be adjusted, within a small range, by mounting one of the mirrors onto a piezoelectric actuator. The stabilization of f_r can be achieved by driving the piezo with the error signal arising from the phase comparison between the detected repetition rate and a stable RF or microwave source. An illustration of the locking set-up is shown in Figure 2.9. In this case the reference source must have the same fractional stability that is wanted for the optical modes. As the actuator is based on a piezoelectric material, the stabilization bandwidth will typically extend only up to a few kHz. Higher bandwidths can be achieved by integrating an electro-optic modulator (EOM) into the cavity (intra-cavity modulator), but this is yet to become a standard feature in commercial systems.

Stabilization to optical references If we exclude the cryogenic sapphire oscillator, the stability of the RF or microwave frequency references found in most laboratories is never better than several parts in 10^{14} for time scales longer than 1 s. This means that the optical frequencies of an optical comb whose repetition rate is phase locked to one of these frequency references will exhibit an average fluctuation in excess of 10 Hz over these time scales. A much better stability can be achieved if the frequency comb is referenced to an optical frequency source rather than an RF or microwave source. This can be accomplished by phase locking one of the optical modes to the optical reference whilst stabilizing f_{CEO} at the same time. As the optical frequency comb has only two degrees of freedom, all the other optical modes will be stabilized as well and hence the spacing between them.

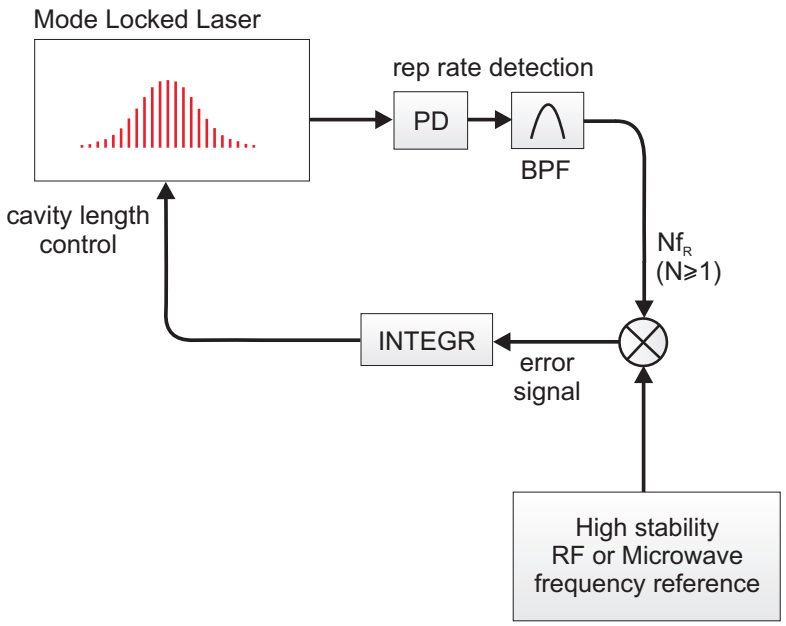

Figure 2.9: Set-up for phase-locking the repetition rate to a high stability RF or microwave frequency reference. The fundamental or a harmonic of the repetition rate is detected and selected by a narrow band-pass filter (BPF). It is then phase compared with the frequency reference by means of a mixer. The error signal, resulting from the phase comparison, is then used to control the cavity length of the mode-locked laser. Integr: integrator; PD:photodetector; BPF: band-pass filter.

The locking of a comb to an optical reference is usually achieved via the "offset locking" technique. In this technique an optical beat is formed between the optical reference and one of the modes of the comb (not necessarily the closest one in frequency). In the case of free-space optical signals the comb and the optical reference beams are combined colinearly and sent to a grating. The grating is used to spatially separate the optical modes of the comb such that only a small number of them (those with a frequency close to that of the optical reference) reach the photodetector. This reduces the shot noise contribution from the modes that do not contribute to the extracted beat frequency. In the case of fibre coupled optical signals the grating is replaced with a narrow band pass optical filter. The beat note is then phase compared to a synthesizer set at the same frequency and the error signal thus obtained is used to control the repetition rate of the frequency comb.

The phase locking scheme is illustrated in Figure 2.10. To understand how the scheme works, I recall here that an optical frequency comb can be fully described by two variables: the repetition rate frequency and the carrier-envelope offset frequency (CEO). Knowing both variables, we can use the comb to measure the frequency of an unknown optical carrier as:

$$f_{\mathrm{opt}} = N f_{\mathrm{r}} + f_{\mathrm{ceo}} + f_{\mathrm{b}} \tag{2.3.5}$$

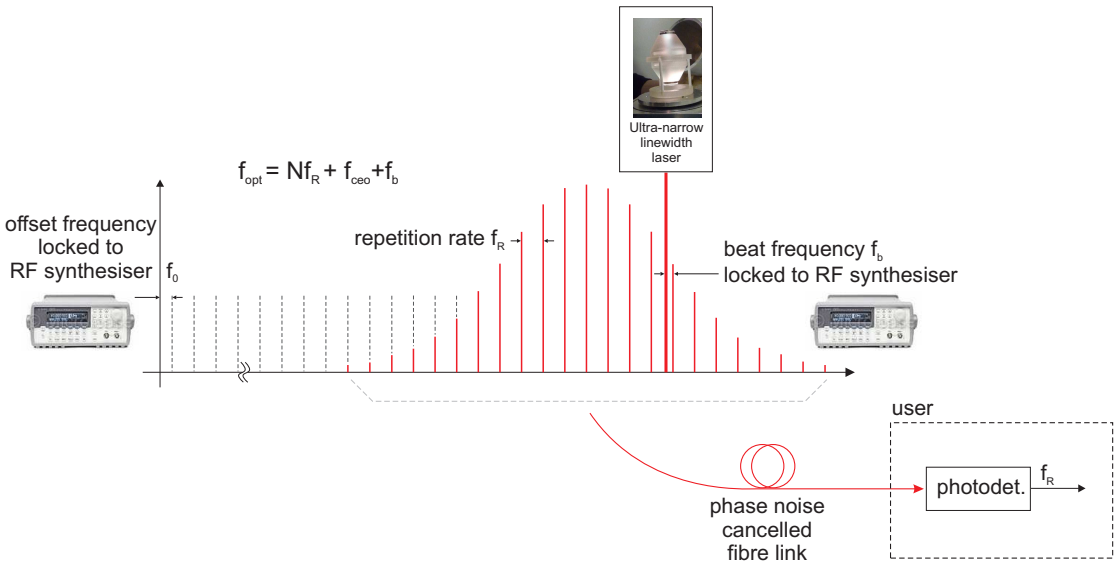

Figure 2.10: Illustration of a mode-locked laser locked to a high stability optical reference. By phase locking one of the teeth of the comb to the optical reference and the carrier-envelope offset frequency to an RF synthesiser, the repetition rate becomes fixed. Any frequency fluctuations of the offset frequency and the beat frequency appear on the repetition rate suppressed by the mode number N.

where f_{r} is the repetition rate frequency, f_{ceo} is the carrier envelope offset frequency and f_{b} is the beat between a comb tooth and the optical carrier to be measured. From the same relationship we can see that if we phase-lock f_{b} and f_{ceo} to external RF references, the repetition rate becomes fixed. Rewriting the equation for f_r

$$f_{\mathrm{r}} = \frac{f_{\mathrm{opt}} - f_{\mathrm{ceo}} - f_{\mathrm{b}}}{N} \tag{2.3.6}$$

one can notice that the frequency stability requirement of both RF references is not demanding as any fluctuation of their frequency causes a deviation of the repetition rate frequency suppressed by a factor of N. For a 1.56 μm mode-locked laser (~200 THz) with 100 MHz repetition rate, N is of the order $\sim 2{\times}10^6$. Typical values of the CEO frequency are in the region of a few tens of MHz and, for fibre lasers, the repetition rate frequency generally can vary between a few tens of MHz to a few hundreds of MHz. The beat frequency f_{b} is, by construction, lower than half of the repetition rate frequency and f_{ceo} and f_{b} are many orders of magnitude smaller than f_{opt}. In the experiments described in this thesis, f_{ceo}=20 MHz and f_{r}=100 MHz. It follows that

$$\frac{f_{\mathrm{opt}}}{f_{\mathrm{ceo}}} \simeq 10^7 \tag{2.3.7}$$

and

$$\frac{f_{\mathrm{opt}}}{f_{\mathrm{b}}} \geq 4 \times 10^{6}. \tag{2.3.8}$$

We can rewrite equation 2.3.6 in fractional terms as

$$\frac{\Delta f_{\mathrm{R}}}{f_{\mathrm{R}}} = \frac{\Delta f_{\mathrm{opt}} - \Delta f_{\mathrm{ceo}} - \Delta f_{\mathrm{b}}}{f_{\mathrm{opt}} - f_{\mathrm{ceo}} - f_{\mathrm{b}}} \tag{2.3.9}$$

where the symbol Δ indicates a frequency change. Remembering that $f_{\mathrm{b}}, f_{\mathrm{ceo}} \ll f_{\mathrm{opt}}$, equation 2.3.10 can be written as

$$\frac{\Delta f_{\mathrm{R}}}{f_{\mathrm{R}}} \simeq \frac{\Delta f_{\mathrm{opt}}}{f_{\mathrm{opt}}} - \frac{\Delta f_{\mathrm{ceo}}}{f_{\mathrm{opt}}} - \frac{\Delta f_{\mathrm{b}}}{f_{\mathrm{opt}}}. \tag{2.3.10}$$

In order to preserve the stability of the optical reference, we must have that both

$$\frac{\Delta f_{\mathrm{ceo}}}{f_{\mathrm{opt}}}, \frac{\Delta f_{\mathrm{b}}}{f_{\mathrm{opt}}} < \frac{\Delta f_{\mathrm{opt}}}{f_{\mathrm{opt}}}. \tag{2.3.11}$$

Taking 10^{-15} as the value of the fractional stability of the optical reference, we can write

$$\frac{\Delta f_{\mathrm{ceo}}}{f_{\mathrm{ceo}}} < \frac{\Delta f_{\mathrm{opt}}}{f_{\mathrm{opt}}} \cdot \frac{f_{\mathrm{opt}}}{f_{\mathrm{ceo}}} = 10^{-8} \tag{2.3.12}$$

and

$$\frac{\Delta f_{\mathrm{b}}}{f_{\mathrm{b}}} < \frac{\Delta f_{\mathrm{opt}}}{f_{\mathrm{opt}}} \cdot \frac{f_{\mathrm{opt}}}{f_{\mathrm{b}}} = 4 \times 10^{-9}. \tag{2.3.13}$$

These values of frequency stability can be readily achieved with ordinary laboratory synthesisers. With this condition verified, the fractional stability of the repetition rate is approximately that of the optical carrier:

$$\frac{\Delta f_{\mathrm{R}}}{f_{\mathrm{R}}} \simeq \frac{\Delta f_{\mathrm{opt}}}{f_{\mathrm{opt}}}. \tag{2.3.14}$$

Equation 2.3.6 is valid only if the division by N can be considered ideal. Research on this matter has shown that the division process can be assumed to be noiseless well beyond the 10^{-15} level (Diddams et al., 2002; Ma et al., 2004, 2007), so for any practical purpose of the comb transfer technique.

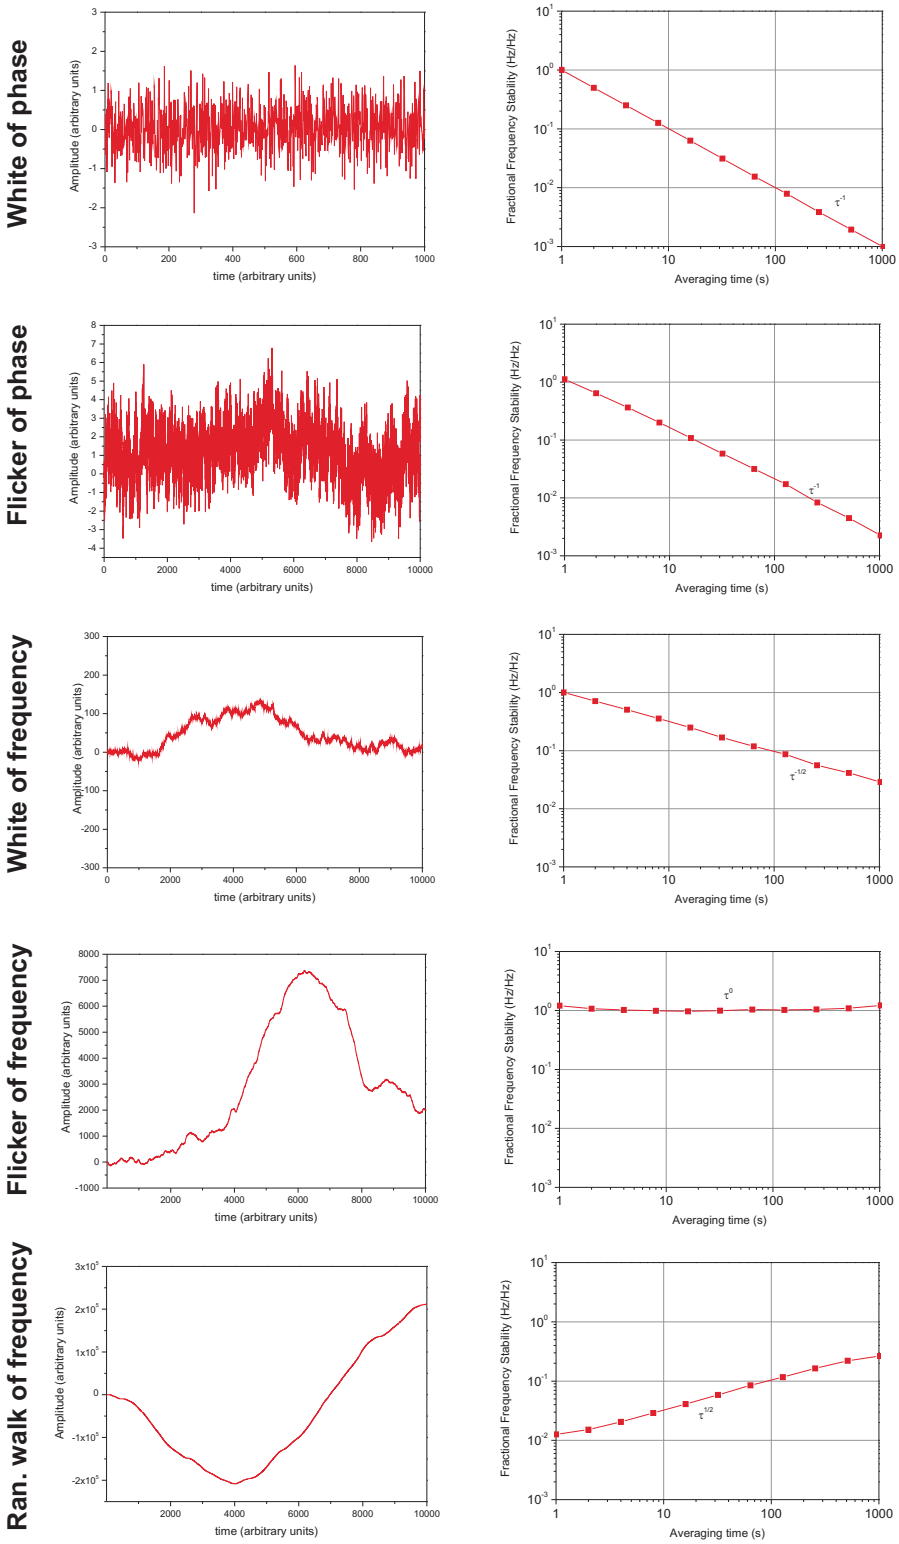

Figure 2.11: Allan deviation slopes for the 5 most commonly found noise processes.

Chapter 3

Techniques for frequency transfer over optical fibre

The research activity on the transfer of ultra-stable frequencies over fibre links originated in the early 1990s and has grown considerably in the last five years. Three techniques have been primarily investigated to date:

1. microwave transfer by intensity modulation of an optical carrier;
2. direct transfer of an optical carrier;
3. microwave and optical frequency transfer with a frequency comb;

In all these techniques the enviromentally-induced fibre noise is suppressed using a two-way technique whose principle is depicted in figure 3.1. A very stable signal of phase ϕ_i is propagated through a fibre to the user end. Acoustic noise and temperature variations cause the optical path length to vary over time, converting into a phase error $\phi_e(t)$ on the signal received at the user end. In order to cancel these phase changes, a portion of the signal is returned to the transmitter end and phase compared with the signal initially injected into the fibre. Under the assumption that the phase changes are much slower than the time that it takes for the light to travel a round trip, we can say that at the transmitter end the observed phase changes are $2\phi_e(t)$. If, by means of an actuator, the phase of the signal before it enters the fibre is corrected by exactly the amount $\phi_e(t)$, the noise at the user end will be suppressed. Considerations on how efficiently this cancellation can be performed will be made in following chapters.

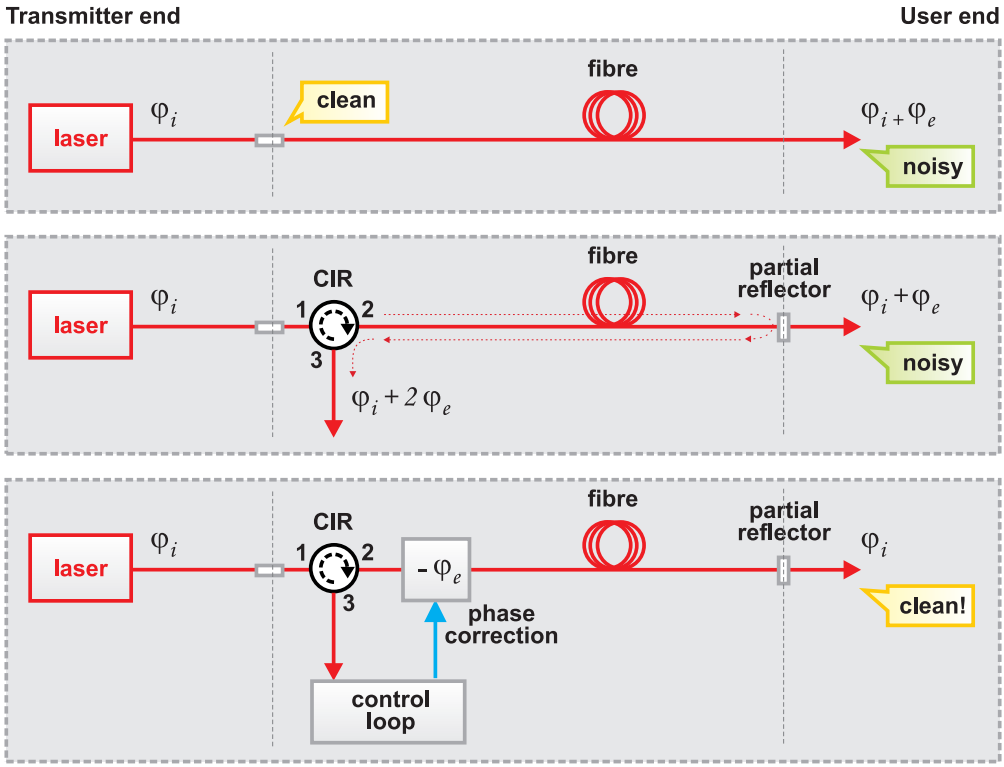

Figure 3.1: Illustration of the principle used for the cancellation of the enviromentally-induced phase noise in a fibre link. CIR: circulator;

3.1 Microwave transfer by intensity modulation of an optical carrier

In this technique a microwave frequency is transferred to the remote end of the fibre by modulating the intensity of a free running laser. The modulation can be imposed either by driving directly the current of the laser or with an external intensity modulator such as a Mach-Zender interferometer. At the far end of the fibre a photodetector retrieves the modulation frequency. An illustration of the microwave transfer over fibre is shown in figure 3.2. The technique has proven to be suitable for the transfer of the best microwave and optically-derived frequency references available today with long term transfer stabilities as high as parts in 10^{-19}.

Preliminary experiments on the transfer of a microwave frequency over a fiber dates from as early as 1980 by the Jet Propulsion Laboratory (JPL) (Lutes, 1980). The work was driven by the stringent synchronization requirements of the Deep Space Network (DSN). The DSN is an international network of antennas placed approximately 120 degrees apart around the world, in California's Mojave Desert (USA), near Madrid (Spain) and Canberra (Australia). The network allows for far-reaching spacecraft to be constantly observed as the Earth rotates. In each of the three locations, multiple antennas need to be synchronized to a master reference (typically a hydrogen maser) that can be up

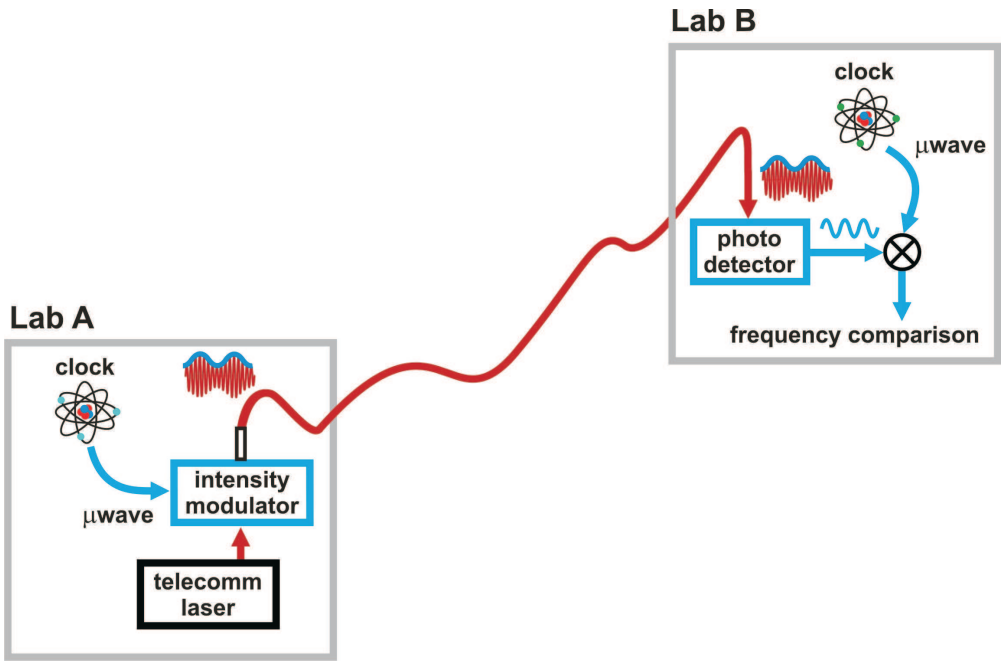

Figure 3.2: Illustration of the microwave transfer over a fibre link principle.

to 30 km away. Whilst in the 70s reference distribution was achieved using microwave coaxial cables, the advent of optical fibres opened the possibility of achieving improved performance and experimental tests on fibre links were initiated. In 1986 a 100 MHz signal from a hydrogen maser was successfully distributed between antennas located 7 km apart (Lutes, 1986) without implementing active compensation of the enviromentally induced optical path fluctuations of the fibre. In 1993, by using a thermally controlled fibre spool to correct for these fluctuations, the JPL scientists showed improvements of the long term transfer stability (1000 s) by more than an order of magnitude to parts in 10^{17} over a 3.8 km-long spooled fibre. A few years later (1994), the JPL group showed results for a microwave, rather than RF, carrier (8.4 GHz) transmitted over a 12-km installed fibre link by externally modulating the intensity of a Nd:YAG laser (Yao et al., 1994). However, no fibre noise cancellation was implemented in this work. In 2002 results for a compensated 16-km spooled fibre link were shown demonstrating a stability of parts in 10^{14} at 1 s and parts in 10^{17} at longer time scales (Calhoun et al., 2002). In this experiment the actuator was a thermally-controlled fibre spool and the fibre link was buried 1.5 m underground. A summary paper of the JPL research on distribution of stable frequency references over fibre can be found in (Calhoun et al., 2007).

A few years later, the technique was taken further by LPL-SYRTE with a series of microwave frequency transfer experiments at progressively higher modulation frequency, 100 MHz (Daussy et al., 2005), 1 GHz (Lopez et al., 2008) and 9.15 GHz (Lopez et al., 2010). The transfer stability measured at 1 s was respectively 1.3×10^{-14}, 5×10^{-15} and 1.3×10^{-15}. The highest stability reported at long time scales was achieved with the 9.15 GHz modulation reaching the impressive level of 10^{-19} for averaging times

Year	Lab	L (km)	$\sigma_x(1s)$	$\sigma_x(1000s)$	MBW (Hz)
(Daussy et al., 2005)	LPL-SYRTE	86	3×10^{-14}	1.3×10^{-16}	10
(Lopez et al., 2008)	LPL-SYRTE	86	5×10^{-15}	4×10^{-17}	3
(Kumagai et al., 2009)	NICT	57	1.5×10^{-15}	n/a	5
(Lopez et al., 2010)	LPL-SYRTE	86	1.3×10^{-15}	7×10^{-18}	15
(He et al., 2011)	NMI(AU)	$20^{(*)}$	1×10^{-13}	5×10^{-16}	15

Table 3.1: Experimental results from different research groups for the transfer of a microwave frequency over installed fibre by intensity modulation of a CW laser. NMI(AU): national measurement institute, Australia; L: length of the fibre link; $\sigma_x(1s)$: fractional stability at 1 s; $\sigma_x(1000s)$: fractional stability at 1000 s; MBW: measurement bandwidth; $^{(*)}$:spooled fibre; Lab: laboratory.

longer than 5×10^4 s. In these experiments a fibre stretcher was used, in addition to a thermally-controlled fibre spool, to achieve higher control bandwidth for the compensation of vibration and temperature-induced phase fluctuations of the transferred microwave signal. Another interesting feature of the LPL-SYRTE experiments was the addition of a polarization scrambler to overcome limitations imposed by polarization mode dispersion. The phase noise cancellation technique can work only if the fibre-induced noise is the same in both directions of propagation of the optical signal over time scales much smaller than the round-trip time. As installed fibre is not usually polarization-mantaining, the light propagates differently during the forward and backward trip. This causes a violation of the required symmetry of the phase noise in the two directions of propagation, limiting the stability of the signal being transmitted.

A different fibre noise cancellation approach, in which the phase noise correction was performed in the microwave rather than optical domain, has been implemented in (Primas et al., 1988, 1989) and more recently in (Zhang et al., 2011). A summary of recent microwave transfer experimental results is shown in Table 3.1.

3.2 Direct transfer of an optical carrier

In this technique, a single optical carrier is directly transferred to the user. In actual experimental setups, an acousto-optic modulator is inserted at the beginning of the fibre link, allowing for the phase of the optical carrier delivered to the user to be corrected in opposition to the phase of the enviromentally-induced phase changes introduced by the fibre. Phase fluctuations slower than the response time of the feedback loop can thus be suppressed at the user end. It should be noticed that, contrary to the intensity-modulated laser microwave technique and the comb transfer technique discussed in chapters 5,6 and 7, no stabilization of the optical path takes place here. What is cancelled is the Doppler-induced frequency shift of the optical carrier whilst the optical path length varies over time.

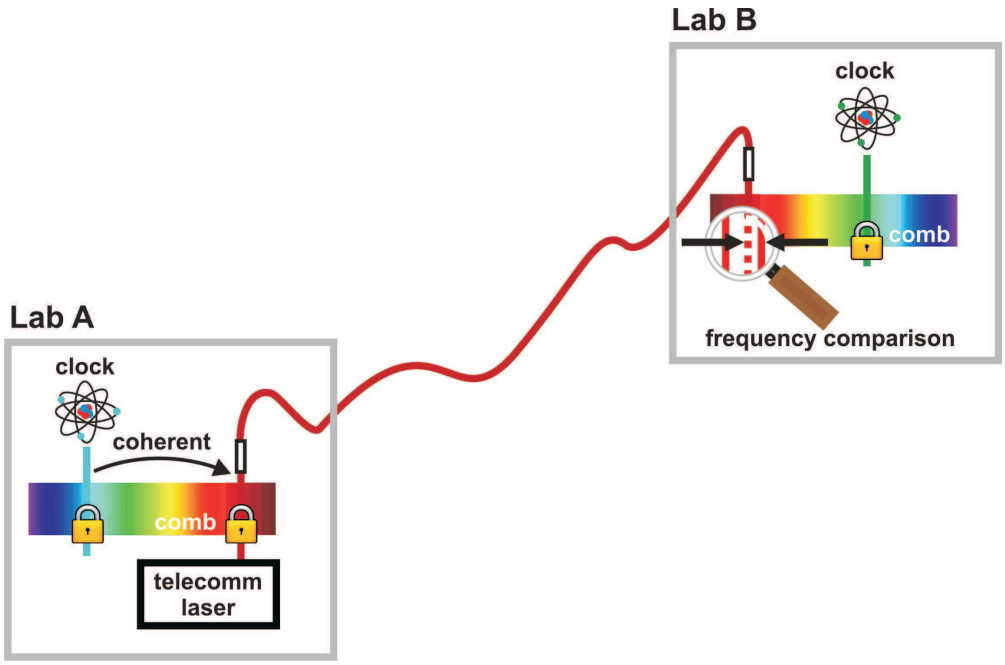

Figure 3.3: Illustration of the optical carrier transfer over a fibre link principle.

Year	Lab	L (km)	$\sigma_x(1s)$	$\sigma_x(1000s)$	MBW (Hz)
(Jiang et al., 2008)	LPL-SYRTE	86	1.3×10^{-16}	5×10^{-19}	10
(Jiang et al., 2008)	LPL-SYRTE	172	4×10^{-16}	1.4×10^{-18}	10
(Williams et al., 2008)	NIST	76	1×10^{-16}	3×10^{-19}	–
(Musha et al., 2008)	NMIJ	110	8×10^{-16}	n/a	-
(Grosche et al., 2009)	PTB	146	3×10^{-15}	3×10^{-18}	$> 10^3$
(Fujieda et al., 2011)	NICT	90	2×10^{-15}	5×10^{-18}	$> 10^3$
(Predehl et al., 2012)	MPQ/PTB	920	4×10^{-14}	3×10^{-17}	–
(Parker et al., 2012)	NPL	118	3×10^{-16}	6×10^{-18}	7

Table 3.2: Experimental results from different research groups for the direct transfer of an optical carrier over installed optical fibre. Lab: laboratory; L: length of the fibre link; $\sigma_x(1s)$: fractional stability at 1 s; $\sigma_x(1000s)$: fractional stability at 1000 s; MBW: measurement bandwidth;

The first experiment demonstrating a high stability transfer of an optical carrier on a noise-cancelled 25 m-long link was performed by Ma (Ma et al., 1994) in 1994 at JILA. The same cancellation principle, described in the previous sections, was then reused a few years later in experiments in number of NMIs around the world over installed fibre lengths up to many hundreds of km. A summary of the transfer stabilities achieved by the different laboratories is shown in Table 3.2.

Beyond demonstrating high levels of stability over installed fibre links longer than 100 km, other interesting tests were undertaken in view of an European (and possibly world-wide) time and frequency dissemination optical network for the comparison of state of the art atomic clocks. The LPL-SYRTE group has demonstrated the transfer of optical

frequencies over urban networks carrying internet traffic (Kéfélian et al., 2009; Olivier et al., 2012). This is an important step towards a many-hundreds of km link. Long-haul dark-fibre links are in fact too expensive to hire and a single ITU channel on a fibre carrying data traffic on the other channels is likely to be the only affordable solution. Cross-talk between the channel used for high speed data transmision and that used for the optical carrier transfer could potentially degrade the transfer stability achievable. However, no degradation was found in the LPL-SYRTE experiments. Another very important step forward towards a realistic link connecting NMIs located in different countries was the demonstration of a 920-km link between PTB and MPQ (Predehl et al., 2012).

Optical repeaters and bidirectional optical amplifiers have been explored by LPL-SYRTE and PTB respectively in view of transfer distances exceeding 1500 km for international optical clock comparison. The optical repeaters are required to overcome the limitations on the noise cancellation bandwidth imposed by the round-trip time when the fibre link exceedes a few hundred km. With a 1000 km-long optical link, for example, the round trip time is 5 ms, limiting the noise cancellation bandwidth to less than 50 Hz. If a repeater is used every 200 km, this bandwidth can be increased by a factor of 5 to 250 Hz. A repeater is essentially a laser that is phase-locked to the stable optical carrier incoming from the preceeding span and and an AOM-based system that performs the fibre noise cancellation of the following span. Bidirectional amplifiers, which are not available commercially, are required along the link to compensate for the fibre loss (which can significantly exceed 200 dB even in real links much shorter than 1000 km) whilst allowing the propagation of the laser beams in both directions, a condition required for the fibre noise cancellation. PTB has also investigated the use of Brillouin amplifiers (Terra et al., 2010). Brillouin amplifiers could be good candidates to reduce the required number of amplifiers in long-haul link as they can amplify optical signals as low as -60 dBm.

3.3 Microwave and optical frequency transfer with a frequency comb

The previous techniques are suitable for the transfer of either a microwave or an optical frequency. Simultaneous transfer of microwave and optical frequencies over fibre can be achieved by propagating an optical frequency comb. At the user end, a large number of harmonically-related microwave frequencies can be extracted with a photodiode or the optical modes can be directly used. The fibre phase noise detection, as will be shown in the following chapters, can be performed either after optical-to-microwave conversion or directly in the optical domain.

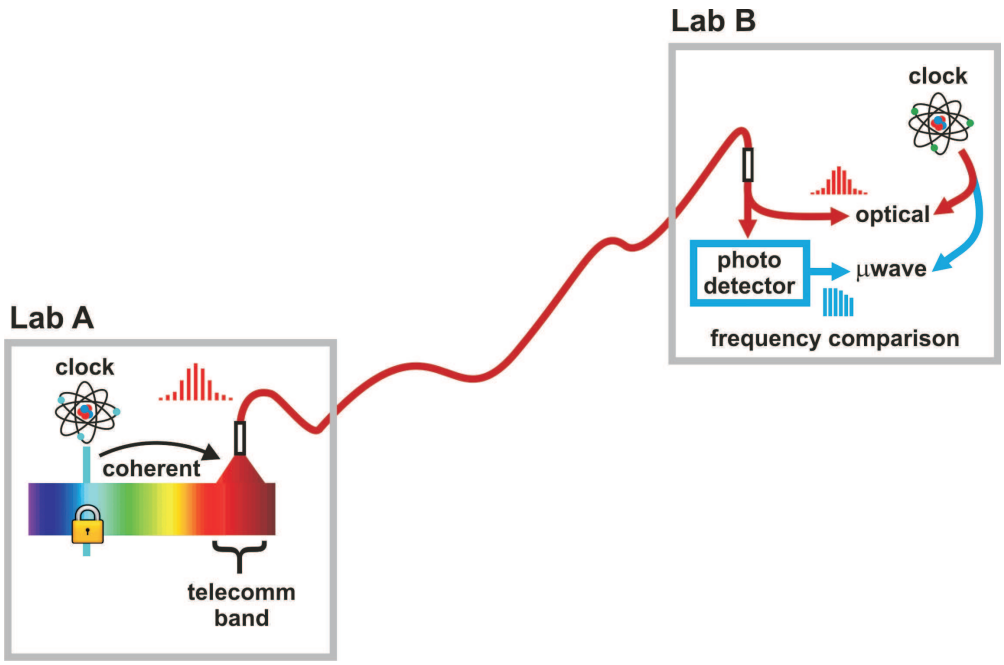

Figure 3.4: Illustration of the principle of the simultaneous microwave and optical frequency transfer using an optical frequency comb.

The earliest work that I am aware of on transfering a frequency using a comb propagated over a fibre was by JILA (Holman et al., 2004). In their first paper in 2004 they transmitted an optical frequency comb of optical bandwidth up to 15 nm over 6.9 km of installed fibre, showing a transfer stability of 3×10^{-14} at 1 s. In this experiment no fibre noise cancellation was implemented and no dispersion compensation was used. When these two additional features were implemented in a later experiment in late 2004 (published in the following year (Holman et al., 2005)), a stability of 9×10^{-15} at 1 s was measured.

Recently, a different phase noise correction approach was used, in which the fibre-induced fluctuations of the transmitted microwave carrier were compensated for by adjusting the laser repetition rate rather than the optical path length of the fibre link (Hou et al., 2011). However, with this technique the measured transferred stability was limited at the level of 10^{-12}.

Experiments were also performed on the transfer of microwave frequencies using optical combs in free space rather than optical fibre (Gollapalli and Duan, 2010). The transfer was performed over a 60 m-long free space optical path achieving a transfer stability of parts in 10^{12} at 1 s and just above 1×10^{-15} at 500 s.

The work that is presented in this thesis was largely inspired by the work pioneered by JILA and taken forward by achiving improved performance at the same time as using longer fibre lengths. Several aspects of comb-based frequency transfer techniques have been investigated in more detail and the experiments have been designed to overcome

some important limitations of JILA experiments, such as the way the actual tranfer stability is measured (this will be discussed in more detail in 5) and the limited measurement time. In the last experiment described in this thesis, the sensitivity of the fibre-noise detection technique has also been improved by orders of magnitude by performing it in the optical rather than microwave domain. The comb transfer technique was also taken forward by measuring for the first time the stability and accuracy of the optical comb modes, rather than just their frequency spacing, for a noise-cancelled link and both were found to be better than a few parts in 10^{18}.

All three techniques discussed in this chapter allow a frequency to be transferred over an optical fibre. However, depending on the application, one particular technique might be preferable. For the comparison of clocks between remote laboratories, certainly the preferred solutions are certainly the direct optical carrier transfer and the intensity-modulated carrier techniques as the optical bandwidth required in these cases is minimal, making the use of commercial networks more affordable as only a single ITU channel is required. The direct optical carrier transfer is to be preferred against the intensity-modulated carrier for its superior performance and reduced complexity. When a dark-fibre network is available, the comb transfer technique provides a convenient way of transferrig both microwave and optical frequencies at high stability and accuracy levels. Pulse trains generated by mode-locked lasers have already found application as timing distribution systems in particle accelerators (Löhl et al., 2010; Kärtner, 2012) and the novel work presented in this thesis will hopefully expand the range of applications of this technique.

Chapter 4

Optical to microwave conversion

Optical to microwave conversion is a crucial process for microwave frequency transfer using optical frequency combs. In this chapter I will discuss how some parameters, such as the width and shape of the propagated pulses, the noise of the laser source and the noise introduced during the detection process can affect the quality of the extracted microwave signals.

4.1 Repetition rate detection

In principle, any harmonic of the repetition rate could be chosen to extract the phase information required for the compensation of the noise introduced by the fibre. Exactly the same information is contained in each harmonic. However, in practice it is more advantageous to extract the highest possible harmonic of the repetition rate, compatible with the availability of high speed photodiodes and high frequency amplifiers. The reasons are the following:

Amplitude noise versus phase noise One important concept to take into consideration is that whilst the phase noise increases as the square of the ratio between the frequencies of the high and low harmonics chosen, the amplitude noise remains the same. When expressed in dB, the noise increase corresponds to $20 \cdot \log(N)$, where N is the frequency ratio between the high and low harmonics. This means that by choosing the highest possible harmonic, we can minimize the contribution of the AM noise. I have measured the AM noise of different harmonics of the repetition rate extracted directly at the output of the laser using a microwave power detector and it is shown in Figure 4.1. The measurement shows that it does not change when different harmonics are selected, as expected. The AM noise is an important parameter in the frequency transfer work described in this thesis as it can easily become indistinguisable from the fibre phase

noise in the detection techniques discussed earlier (mainly because of the residual AM sensitivity of microwave phase detectors), leading to unsuppressed noise at the user end.

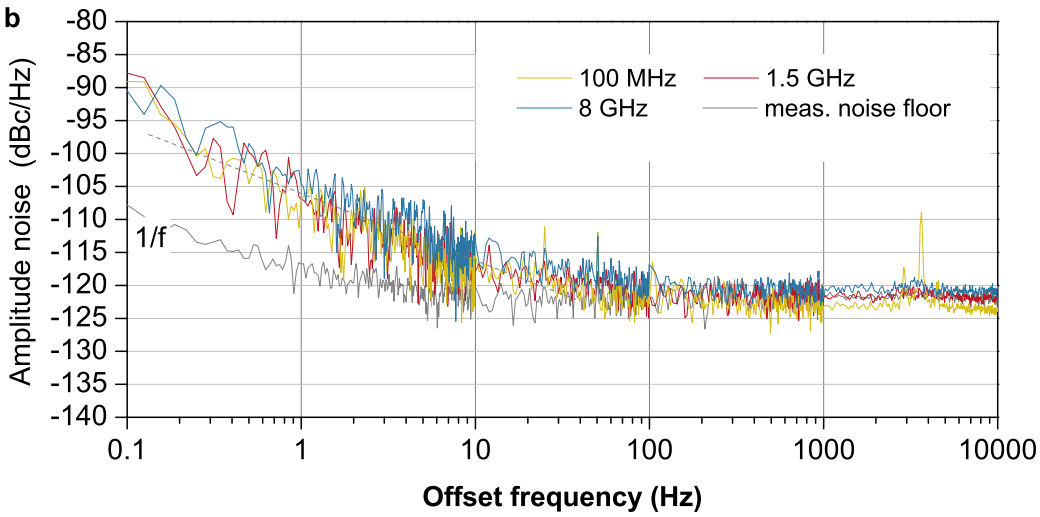

Figure 4.1: AM noise for the fundamental, the 15^{th} and 80^{th} harmonics of the repetition rate.

Shot noise The same signal-to-noise considerations discussed in the previous paragraph are valid for shot noise. Also this type of noise is frequency independent and therefore choosing a higher harmonic is preferrable.

Noise in amplifiers Another reason for choosing the highest possible harmonics of the repetition rate is the intrinsic noise of electronic devices. The additive noise of RF or microwave amplifiers depends primarily on the operating temperature and the details of the technology used, but not on the operating frequency (to first approximation). The output white phase noise depends on the input power and can be calculated as $L_\phi(f) =$ –177 dBc – (signal power in dB) – (NF in dB), where NF is the noise figure and –177 dBc is the equivalent SSB white phase noise due to additive Johnson noise ($kT_0 = 4\times10^{-21}J$) at the temperature of 290 K (RubiolaTutorial). Experimentally one can observe that, with a repetition rate of 100 MHz and with regular photodiodes, the power available at each extracted harmonic of the repetition rate typically cannot be much higher than –20 dBm owing to saturation effects taking place in the photodiode. This means that at best $S_\phi(f) = -177$ dBm $-(-20$ dBm$) = -157$ dBc/Hz in the case of an ideal amplifier[1] (noise figure = 0 dB). This level of noise corresponds to approximately 1.7×10^{-14} at 1 s at the frequency of 100 MHz (fundamental of the repetition rate) when a 100 kHz

[1]The calculated level is realistic only if the amplitude noise of the mode-locked laser is sufficiently low. In pratical cases, saturation effects in the photodiodes will convert the amplitude noise of the laser into phase noise, making it particularly hard to achieve the calculated level of phase noise.

measurement bandwidth is used. However, if the 10 GHz harmonic is extracted, the same level of phase noise corresponds to a fractional frequency stability of 1.7×10^{-16} at 1 s, two orders of magnitude better than that achievable with the fundamental. One could choose even higher harmonics, but in practice the frequency is generally limited to less than 12 GHz, mainly for cost issues. Although photodetectors and microwave components do exist at frequencies well above 12 GHz, their cost grows exponentially above this threshold.

4.2 Detection of high-order harmonics of the repetition rate

For the reasons discussed earlier, it is desirable to detect the highest possible harmonic of the repetition rate, given the bandwidth of the photodetector and the electronics used to condition the extracted signal. However, the power and signal-to-noise ratio of the extracted harmonics do not just depend on the physical parameters of the optoelectronics involved in the optical-to-microwave conversion, but also on a few parameters of the optical signal itself, as will be shown with the following simple experiments. I ran these experiments at NPL together with Radan Slavík from the ORC, who also provided the specialist equipment required that wasn't available at NPL.

The experimental set-up is shown in Figure 4.2. The pulse train from a mode-locked laser is sent to a wave shaper (Finisar 4000s). In its simplest application, this device can be used as a bandwidth-tunable band pass filter. When used in conjunction with a frequency comb generated by a mode-locked laser, the width of the pulses can be changed by limiting the optical bandwidth over which the comb extends. In a more advanced mode of operation, the waveshaper allows for some degree of dispersion to be introduced to simulate the pulse broadening effect that would be caused by a real optical fibre. Both the filtering and the chromatic dispersion are achieved by adjusting the phase and amplitude of the individual pixels in a liquid crystal-based pixel array.

An EDFA amplifier, set to automatic power control mode, is added at the output of the wave shaper in order to maintain a constant average power delivered to the subsequent optoelectronics. Mantaining a constant optical power onto the photodetector simplifies the analysis as the shot noise levels and the noise introduced by the microwave amplifiers after the photodetector remain the same when the optical bandwidth (and thus the average power) is changed. A bandwidth-tunable optical band-pass filter (Alnair BVF-200) is also used after the EDFA to limit the bandwidth of the amplified stimulated emission (ASE) generated by the amplifier which could be detrimental to the signal-to-noise ratio of the extracted harmonics of the repetition rate. The filter is adjusted to achieve an optical passband always slightly larger than that set with the waveshaper. The filtered optical signal is then split into two outputs: one connected to a high speed photodetector

for the spectral analysis of the harmonics of the repetition rate; the other connected to an optical sampling oscilloscope [2] (Exfo PSO-102) and an optical spectrum analyser for the analysis of the optical pulse train in the time and frequency domain. The measured pulse duration as a function of the optical bandwidth set with the waveshaper is shown in Figure 4.1. The pulse intensity has been normalized to facilitate the comparison of the width of the pulse.

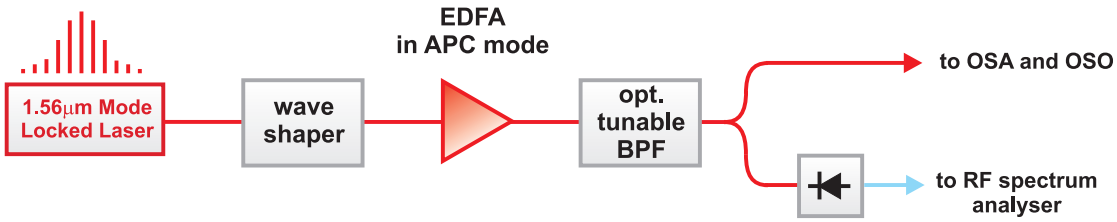

Figure 4.2: Experimental set-up for the measurement of the extracted repetition rate as a function of the duration and the shape of the optical pulses. The EDFA is used in automatic power control (APC) mode. OSA: optical spectrum analyser; OSO: optical sampling oscilloscope.

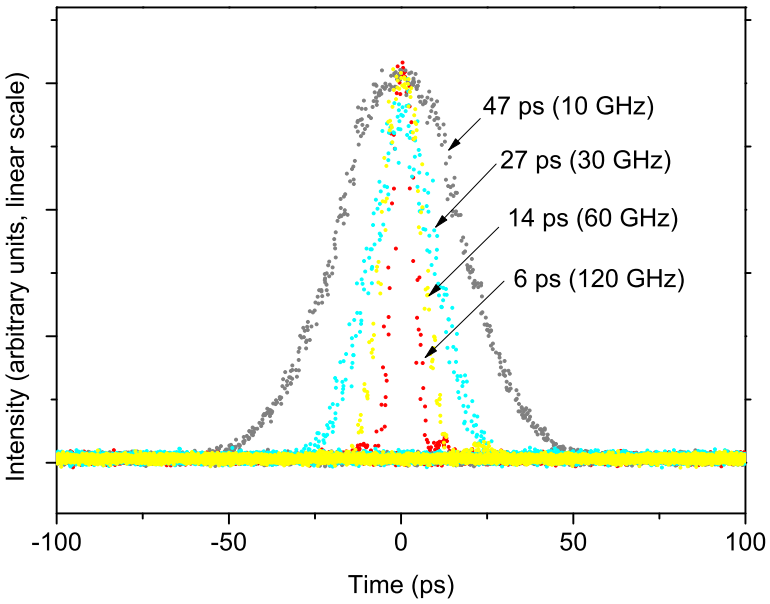

Figure 4.3: Pulse width as a function of the optical bandwidth. The height of the pulses has been normalized for ease of comparison. The maximum pulse width achieved, without dispersion, was 47 ps when the wave shaper filtering bandwidth was set to its minumum allowed value (10 GHz).

The first experiment consisted of measuring the power and the signal-to-noise ratio of the harmonics of the repetition rate as the pulse width is changed. Two harmonics, 1.5 and

[2] An optical sampling oscilloscope uses narrow sampling pulses to open a sampling gate that generates a time-stretched version of the measured signal. Such downconverted signal is then detected by low-speed electronics and digitally processed. The operating principle is the same used in conventional electronic sampling oscilloscopes.

8 GHz, were analysed. The choice of these specific harmonics was mainly driven by the availability of the narrow microwave bandpass filters needed for harmonic selection. The smallest optical bandwidth achievable with the waveshaper is 10 GHz (flat-top shape), corresponding to a transform-limited pulse width of 44 ps (full-width-half-maximum).

In order to evaluate the effect of changes in the pulse duration on the extracted harmonics, we can start by calculating the rise time τ corresponding to these frequencies, which can be approximated with a known formula for a single-pole frequency response

$$\tau_r \simeq \frac{0.35}{f_n}, \tag{4.2.1}$$

where f_n is the frequency of the selected harmonic n. For the 8 GHz harmonic $\tau_r \simeq 44$ ps and for the 1.5 GHz harmonic τ increases to 233 ps. When the rise time of the pulse is higher than that corresponding to the extracted harmonic, changes of the pulse duration should not be detectable in terms of the extracted microwave power. Table 4.1 shows the power and the SNR of the detected harmonics of the repetition rate as the pulse width is increased from 8.5 ps to 47 ps. The RF power and SNR are measured using a spectrum analyser.

BW_{opt} (GHz)	τ_p (ps)	P_{15} (dBm)	SNR_{15}	P_{80} (dBm)	P_{opt} (dBm)
120	8.5	–24	80	–37	–10
60	16	–24	80	–37	-13
30	26	–24	80	–37	–16
10	47	–24	77	–38	–23

Table 4.1: Power and SNR of the 15^{th} GHz and the 80^{th} harmonics of the repetition rate for different optical pulse widths. BW_{opt}: optical bandwidth; τ_p: pulse width; P_{15} power of the 15^{th} harmonic of the repetition rate; SNR_{15}: power of the 15^{th} harmonic; P_{80} power of the 80^{th} harmonic; P_{opt}: power at the input of the optical amplifier

The results confirm the expectation, showing no substantial change of the measured parameters over the allowed pulse width range. The reduction in the SNR of the 1.5 GHz harmonic when the wave shaper is set to the minimum available bandwidth of 10 GHz could be due to increased optical noise of the EDFA as in this case it is driven with an optical power close to its minimum allowed input power level (–25dBm).

In the second part of the experiment we introduced chromatic dispersion with the wave shaper or with some spooled fibre. Using an optical bandwidth of 315 GHz (2.5 nm), the maximum dispersion-induced pulse duration that was possible to achieve with the waveshaper was approximately 50 ps, corresponding to approximately 1.2 km of SMF-28 fibre. When a 4.7 km-long fibre spool was introduced in place of the waveshaper,

the pulse duration increased to 180 ps and 350 ps for optical bandwidths of 315 GHz and 600 GHz respectively. These pulse duration values agree reasonably with those calculated using the dispersion of 17 ps $nm^{-1}km^{-1}$ specified by the manufacturer for a standard SMF-28 fibre from Corning Inc. at 1560 nm. With 315 GHz optical bandwidth we have 17 ps $nm^{-1}km^{-1} \times 2.5$ nm $\times$ 4.7 km $\approx$ 200 ps and with 600 GHz optical bandwidth we have 17 ps $nm^{-1}km^{-1} \times 4.7$ nm $\times$ 4.7 km $\approx$ 380 ps. Figure 4.4 shows the pulse duration and shape for the various conditions, as acquired with an optical sampling oscilloscope. When dispersion is introduced with the wave shaper or the spooled fibre, the pulses assume a square-like shape which maps the flat-top optical spectrum as filtered by the waveshaper (Figure 4.5). The RF powers measured for the above conditions are shown in Table 4.2.

BW_{opt} (GHz)	τ_p (ps)	P_{15} (dBm)	SNR_{15}	P_{80} (dBm)	P_{opt} (dBm)
315	2	–25	80	–37	–5.9
315	50	–25	80	–37	–8.2
315	180	–25	80	–40	–7.0
600	350	–28	74	–55	–5.5

Table 4.2: Power and SNR of the 15^{th} GHz and the 80^{th} harmonics of the repetition rate when dispersion is added. BW_{opt}: optical bandwidth; τ_p: pulse width; P_{15}: power of the 15^{th} harmonic of the repetition rate; SNR_{15}: signal-tonoise ratio of the $15^{th}harmonic$;P_{80}: power of the 80^{th} harmonic; P_{opt}: power at the input of the optical amplifier.

The reduction of the detected power of the 1.5 GHz and 8 GHz harmonics is primarily due to the fact that the pulse is square rather than Gaussian or Lorentian. The Fourier transform of a square pulse train is in fact a $\text{sinc}(x)$ function with nulls at multiples of the inverse of the pulse width. The spectrum of a pulse train of square pulses, like those shown in Figure 4.4, will therefore be modulated in amplitude, as confirmed experimentally in Figure 4.6. The spectrum shown is for a 180 ps-long square pulse train and its minima are found at multiples of 5.4 GHz, in reasonable agreement with the calculated value 1/180 ps = 5.56 GHz. In actual experiments, when the output of the laser is not filtered, the modulation of the RF spectrum is much less pronunced as the spectrum is no longer flat-top. In most cases, the power available at 8 GHz is usually only a few dBs lower than that at 1.5 GHz and at the fundamental.

4.3 Dispersion compensation

In the frequency transfer experiments that will be described in this thesis, a 1.56 μm mode-locked erbium-doped fibre laser emits pulses of duration less than 150 fs with an optical bandwidth extending up to 90 nm (30 nm when optical amplifiers are used), as shown in Figure 4.7. With such a large bandwidth, the pulse-broadening effects due

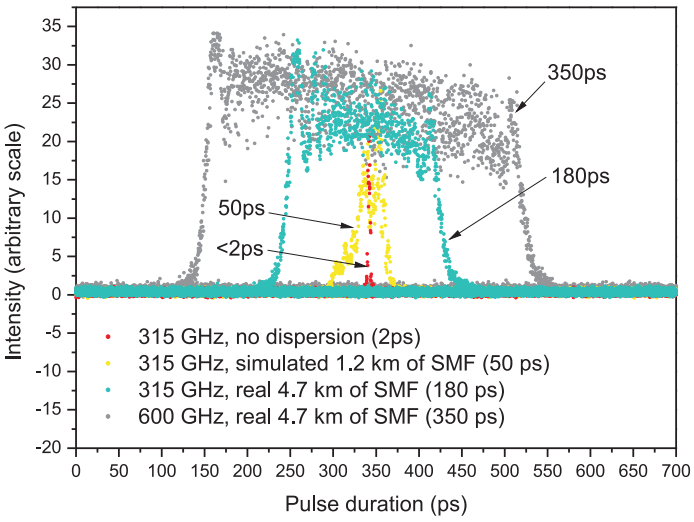

Figure 4.4: Pulse shapes at the input of the photodetector for various optical bandwidth and dispersion values.

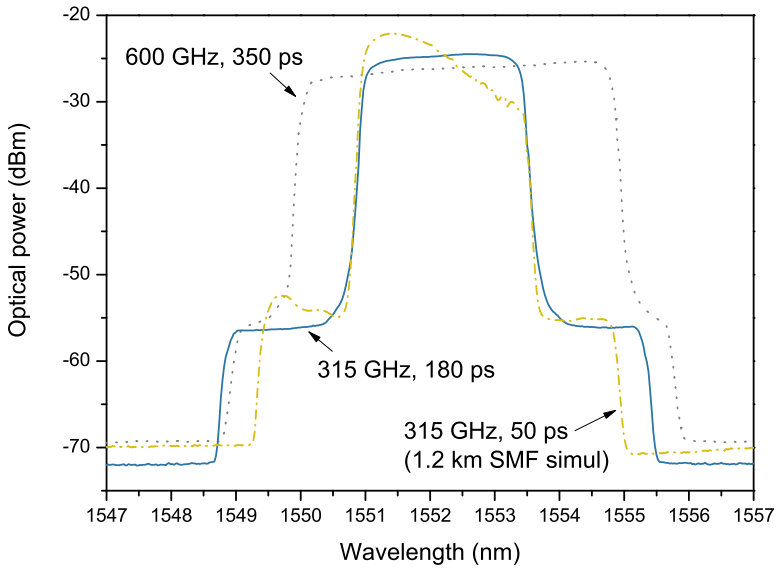

Figure 4.5: Optical spectra of the optical comb as filtered by the waveshaper and the bandwidth-adjustable optical filter.

to the chromatic dispersion of the single mode fibre (SMF) used in the experiment are substantial and dispersion compensation is required in order to recover the repetition rate.

Because of the dispersion of 17 ps/km/nm, the pulses will be broadened to 17 ps/km/nm$\times$ 90 nm $\times$ 100 km = 153 ns after 100 km of fibre. As the period of the repetition rate of the laser used in this experiment is either 4 or 10 ns, the overlap between consecutive

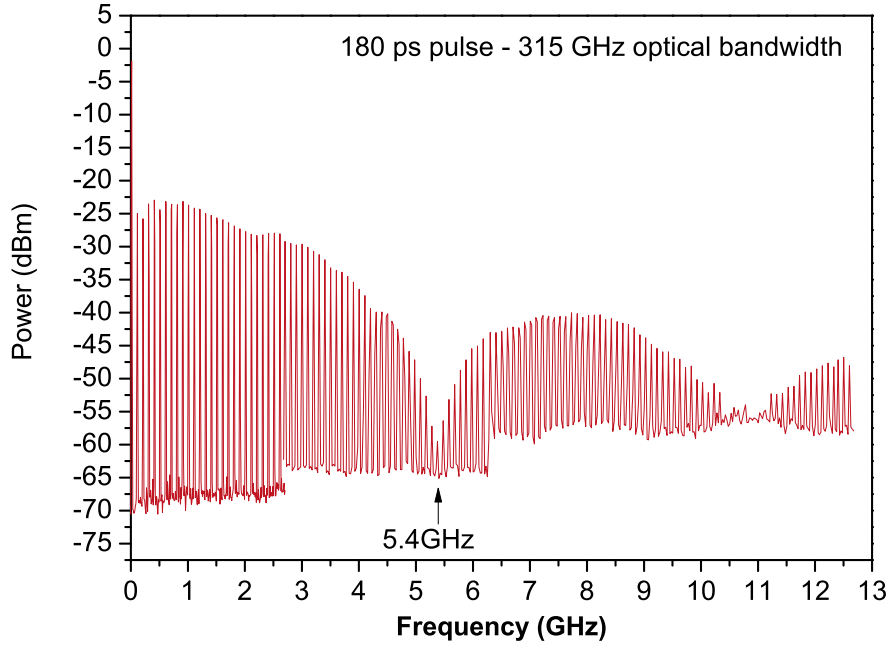

Figure 4.6: Dips in the RF spectrum of the detected harmonics of the repetition rate due to the square shape of the pulses. The frequency of the dips falls approximately at multiple of the inverse of the pulse duration, 180 ps in this case.

optical pulses is considerable. Thus, the laser light at the far end of the fibre will be almost constant rather than pulsed. The dispersion can be greatly reduced by the insertion of a dispersion compensating fibre (DCF) matched to the length of the SMF. With this arrangement, the recovered pulse duration was recompressed typically to less than 200 ps, as illustrated in Figure 4.8. Although the recovered pulse duration is still orders of magnitude worse than that at the output of the laser, as we have seen, it does not compromise the signal-to-noise ratio significantly for harmonics of the repetition rate up to many GHz. Also, it should be noted that the shape of the pulse does not affect the achievable stability of the transferred microwave frequencies (beyond the SNR considerations above) as long as it remains the same over time.

4.4 Self heterodyne effect

In all the experiments that will be described in this thesis, in order to measure the noise of the fibre at the user end and returned at the transmitter end after a round trip, the delivered signal is compared with that initially injected at the 'transmitter' end[3]. In this way, it is possible to measure stabilities far exceeding that of the laser source as

[3]A different approach, that could be viable in principle, would be to use an ultrastable source as the reference which the returned and remote signals are compared against. In practice however, no microwave or optical references can currently meet the stabilities that will be demonstrated here

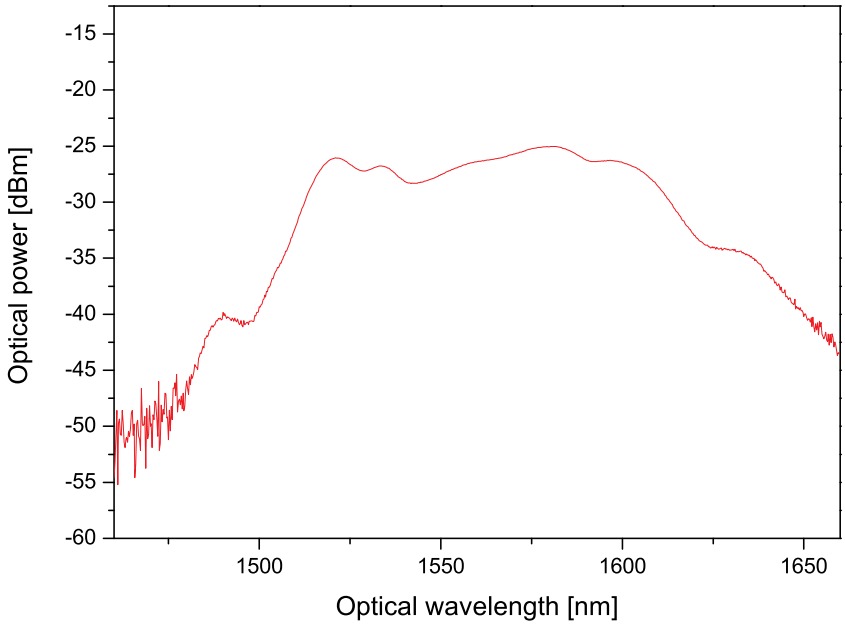

Figure 4.7: Optical spectrum of the mode-locked laser output.

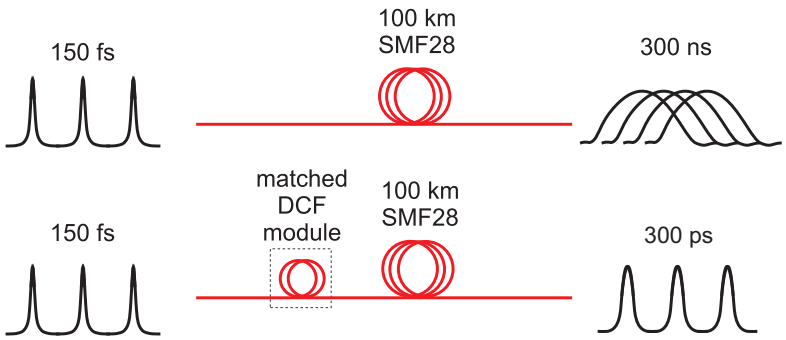

Figure 4.8: The insertion of a matched length of dispersion compensating fibre allows the repetition rate to be recovered after 100 km of fibre.

the latter is 'common' to the signals being measured and therefore rejected. However, as I will show here, the laser source can be considered reasonably 'common mode' only for phase changes much slower than the round trip time of the laser beam in the fibre. The fibre link is effectively a delay line which can cause the noise of the laser source to be revealed at the remote end and after a round trip via the self-heterodyne effect. In Figure 4.9, the set-up for the measurement of the noise introduced by the fibre is shown and the delay is characterized in terms of the Laplace transform.

In the Laplace domain, we can write the output phase fluctuations as a function of the input phase fluctuations via the transfer function $H(s)$

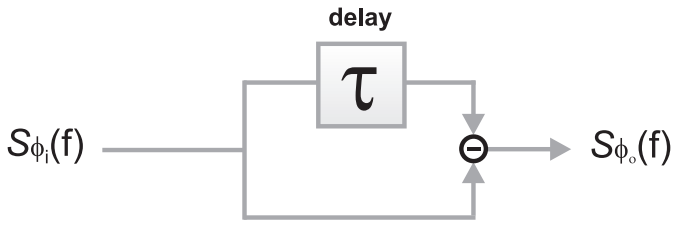

Figure 4.9: Model of the self-heterodyne effect. The optical fibre link introduces a delay τ (proportional to its length) decorrelating the noise of the signals at the input of the phase comparator (represented by a minus sign in this diagram). The phase comparator sees two almost uncorrelated signals where the degree of uncorrelation depends on the ratio between the time scale under consideration and the fibre delay τ. The self-heterodyne-induced noise become more and more a concern as the fibre length is increased and τ increased. However, if ultra stable references (such as ULE-cavity stabilized lasers) are used to stabilized the repetition rate of the mode-locked laser, the effect becomes less important as discussed in the text.

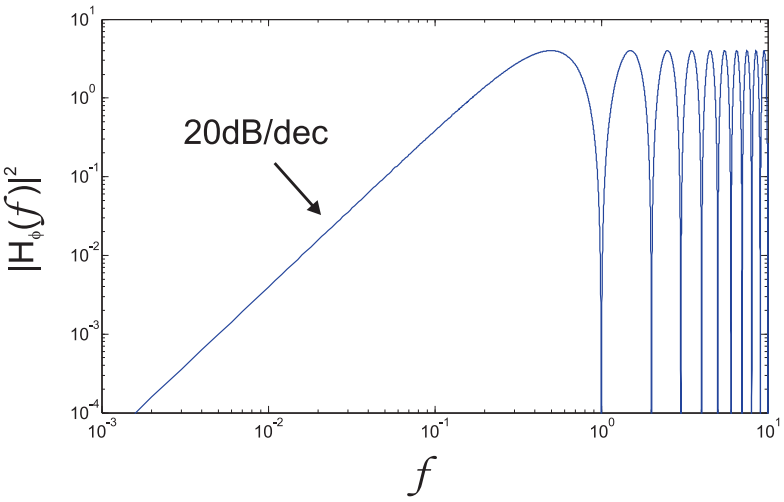

Figure 4.10: Plot of the transfer function $|H_\phi(f)|^2$ that relates the phase noise detected at the output of a delay line, when compared to the input, as a function of the phase noise of the source. The function shows dips at integer values of $1/\tau$ where τ is the delay imposed by the fibre. For offset frequencies well below $1/2\tau$ the noise of the source is suppressed by approximately 20 dB/dec.

$$\Phi_{\mathrm{o}}(s) = H_\phi(s)\Phi_{\mathrm{i}}(s) \tag{4.4.1}$$

where the subscripts o and i stand for *output* and *input* respectively.

In the time domain, the fibre link delays the input phase fluctuations by the amount τ such that $\phi(t) = \phi(t-\tau)$. Using the time-shift theorem of the Fourier transform, which

states that a time shift τ corresponds in the frequency domain to a multiplication of the spectrum by $e^{-s\tau}$, we can write the transfer function corresponding to the measurement setup of Figure 4.9 as follows:

$$H_\phi(s) = \frac{\Phi_o(s)}{\Phi_i(s)} = \frac{\Phi_i(s) - \Phi_i(s)e^{-s\tau}}{\Phi_i(s)} = 1 - e^{-s\tau} \tag{4.4.2}$$

from which, using trigonometric identities, we can write the magnitude of the transfer function

$$|H_\phi(f)|^2 = 4\sin^2(\pi f\tau). \tag{4.4.3}$$

The power spectral density of $\Phi_o(s)$ can be written as

$$S_{\phi_o}(f) = |H_\phi(f)|^2 \, S_{\phi_i}(f). \tag{4.4.4}$$

A plot of $|H_\phi(f)|^2$ is shown in Figure 4.10. The self-heterodyne noise of the mode-locked fibre laser used for the experiments described in this thesis is shown in Figure 4.11, when the repetition rate is locked to a hydrogen maser. As the locking bandwidth is increased and becomes comparable to the first resonance of the piezo actuator (a few kHz) used for the control of the repetition rate in the fibre laser, the noise is enhanced in the 200 Hz to 10 kHz region and the bumps and dips due to the self-heterodyne effect are revealed. The length of fibre used for this experiment was 50 km followed by a dispersion compensating fibre approximately 7 km long. The corresponding delay imposed by these two fibres therefore is $\tau = 5.7 \times 10^4$ km / $2 \times 10^8 \mathrm{ms}^{-1} = 285$ μs, making the frequency of the dips fall at multiples of $1/\tau \approx 3.5$ kHz, which is in very good agreement with the experimental data.

The transfer function $H(jf)$ allows the noise contribution of the reference used to phase lock the repetition rate of the mode-locked laser to be evaluated. In particular, it is interesting to evaluate the self-heterodyne contribution when a ULE cavity-stabilized laser is used as the reference instead of the H-maser and this is shown in Figure 4.11. To calculate the values of the red trace, I have measured the phase noise of two 674 nm ULE cavity-stabilized lasers used for the strontium ion clocks at NPL. The figure also shows the self-heterodyne noise induced when a hydrogen maser is used as the reference. The difference in the noise level between the two cases is of the order of 50 to 60 dB, with the noise contribution from the ULE cavity-stabilized lasers falling below the typical shot noise levels achievable in practical cases. In all the experiments described in this

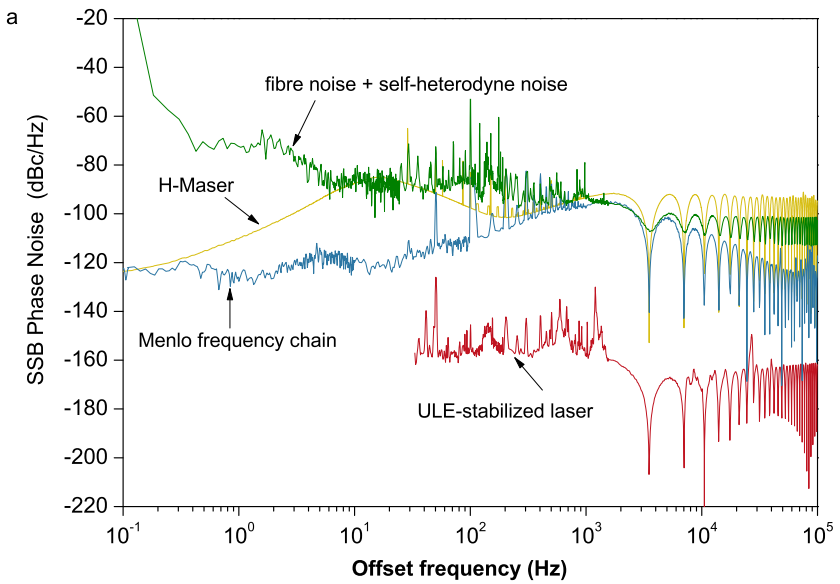

Figure 4.11: Self-heterodyne phase noise detected with a 50 km spool for various frequency references used to lock the repetition rate. Green: maser-referenced synthesizer; Red: calculated noise if a ULE-cavity stabilized laser was used to stabilize the repetition rate; Blue: residual noise of the frequency chain used in the Menlo system to lock the 10th harmonic of the repetition rate; Yellow: phase noise of the maser, filtered with the self-heterodyne transfer function. All values are scaled up to 8 GHz.

thesis the repetition rate was locked to a hydrogen maser as ultra stable lasers were not available continuously. We can therefore conclude that the results that will be shown here could in principle be further improved if an ultra stable laser was used in place of the hydrogen maser. In Figure 4.11 I have also shown the contribution to the detected noise from the frequency chain installed in the Menlo frequency comb. This data tells us that, if a better microwave source was used instead of a hydrogen maser (for example a cryogenic sapphire oscillator), the current frequency chain would need to be improved as it would otherwise limit the achievable noise performance at offset frequencies higher than a few hundred Hz.

4.4.1 Photodetection noise

The conversion of optical pulses into electrical pulses achieved with a photodetector can be affected by the following noise types:

- shot noise;
- thermal noise;

- amplitude to phase conversion.

The shot noise of a photodetector depends on the average DC photocurrent and therefore on the average optical power applied. The power spectral density of the noise produced is independent of frequency up to the maximum bandwidth of the photodetector. The shot noise current is

$$i_{\text{shot}} = \sqrt{2ei_{\text{avg}}} \tag{4.4.5}$$

where e is the charge of the electron and i_{avg} is the average DC photocurrent. When the photodetector is terminated with a load resistance R, the single sideband phase noise due to shot noise can be calculated as

$$\mathcal{L}(f)_{\text{shot}} = \frac{ei_{\text{avg}}R}{P_{\text{signal}}} \tag{4.4.6}$$

where P_{signal} is the power of the RF or microwave signal. In the first experiment that will be described in this thesis, the photocurrent measured after 100 km of fibre (round trip), with 60 mW optical input power, was $\sim 50\ \mu$A and the recovered microwave signal was ~ 100 nW. Using equation (4.4.6) and assuming a 50 Ω load resistance

$$\mathcal{L}(f)_{\text{shot}} = 10\log\left(\frac{ei_{\text{avg}}R}{P_{\text{signal}}}\right) = -144\,\frac{\text{dBc}}{\text{Hz}}. \tag{4.4.7}$$

In later experiments, the optical power incident on the photodetector was usually higher as optical amplifiers were employed. Another source of noise in the photodetection process is the thermal noise of the load resistor required to convert the photocurrent into voltage. The phase noise arising from the thermal noise is

$$\mathcal{L}(f)_{thermal} = \frac{kTR}{2V_{RF}^2} \tag{4.4.8}$$

where k is the Boltzmann constant, T is the resistor temperature and V_{RF} is the rms voltage level of the RF signal. For $R = 50\ \Omega$, $T = 300$ K and a typical value of $V_{\text{RF}} = 2.2\ \text{mV}_{\text{rms}}$

$$\mathcal{L}(f)_{\text{thermal}} = -137\,\text{dBc/Hz}. \tag{4.4.9}$$

Both the shot and the thermal noise are dependent on the optical input power P_{opt}, but while $\mathcal{L}(f)_{\text{shot}}$ scales as the inverse of the optical power, $\mathcal{L}(f)_{\text{thermal}}$ scales as the inverse squared:

$$\mathcal{L}(f)_{\text{shot}} = \frac{ei_{\text{avg}}R}{P_{\text{signal}}} \propto \frac{P_{\text{opt}}}{P_{\text{RF}}} \propto \frac{1}{P_{\text{opt}}}, \tag{4.4.10}$$

$$\mathcal{L}(f)_{\text{thermal}} = \frac{kTR}{2V_{RF}^2} \propto \frac{1}{P_{\text{RF}}} \propto \frac{1}{P_{\text{opt}}^2}. \tag{4.4.11}$$

One noise type will dominate over the other depending on the optical power level applied to the photodetector, as shown in Figure 4.12.

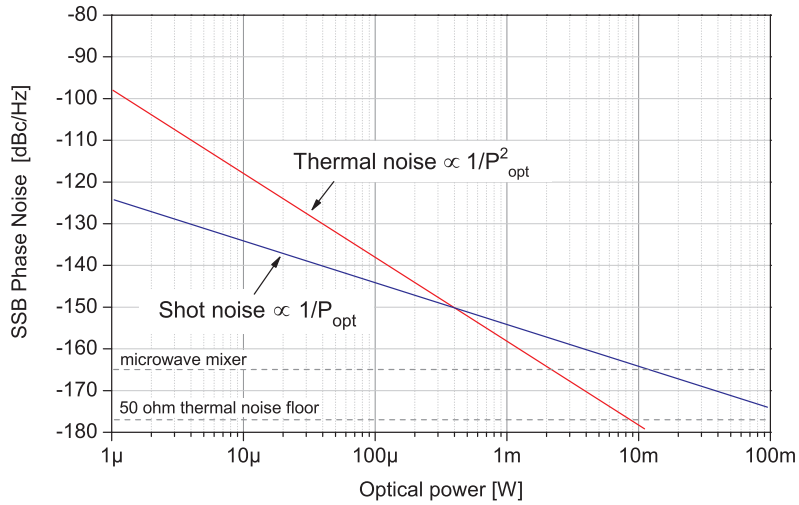

Figure 4.12: Shot and thermal noise as a function of the optical power incident on the photodetector. Both shot and thermal noise decrease with increasing optical power. Above a certain optical power, the noise of other components following the photodetector in the microwave extraction circuit will dominate. The typical white phase noise of a microwave mixer is shown as an example.

One could be tempted to think that high optical power levels are always preferred. However, for high power levels, other effects might occur like amplitude to phase noise conversion (Ivanov et al., 2003, 2005, 2007). For long fibre links, because of the optical

attenuation, it is unlikely that high optical power levels will be available at the user end or after a round trip unless a series of optical amplifiers are used. The consideration of amplitude to phase noise conversion could become more important for the detection of a 'clean' repetition rate signal directly from the laser at the transmitter end, where high optical power is available. However, in all the experiments described in this thesis, the optical power of the repetition rate detected at the output of the laser was kept at levels well below those required to saturate the photodetectors.

4.5 Other sources of noise

In the previous sections we have seen how the noise of the laser source can reveal itself at the user end and after a round trip, affecting the transfer stability that can be achieved. Another important analysis that must be undertaken is that of the noise contributions of the measurement setup. For measurement setup I include not only that for the measurement of the repetition rate delivered at the user end, but also that used for the extraction of the fibre noise at at the local end after a round trip. In the following sections I will show measurements of noise introduced by frequency mixers (used for phase detection), microwave amplifiers and DC amplifiers. The measurements will show at what level the noise of these devices affects the transferred stability. I will also show some measurements of the resolution of frequency counters which are used for the measurement of the transfer stability in a few of the experiments described in this thesis.

4.5.1 Intrinsic noise in microwave mixers

Microwave mixers will be used extensively in the experiments described in this thesis. They will be most often used as phase detectors with the two inputs in "quadrature" as discussed in earlier sections. The noise introduced by a microwave mixer can be measured by feeding both input ports with the same signal as shown in Figure 4.13. Any phase fluctuations of the source are common to both inputs and any resulting voltage at the mixer output will therefore be due only to the mixer itself. Despite the simple set-up, great care is needed for this measurement to achieve sufficient confidence in the result. The measurement can in fact very easily be compromised by the amplitude noise of the microwave source. The microwave mixer when operated as a phase detector has a non-zero sensitivity to the input power level and the source amplitude noise is converted into voltage noise at the mixer output. This effect can be mitigated by choosing a suitable microwave source and by operating the mixer, not in exact quadrature, but at a specific phase difference value that suppresses the amplitude sensitivity. The point at which the AM sensitivity is minimized is often referred to as the "sweet point". The sweet point can be found by intentionally modulating the amplitude of the source and adjusting the phase shift between the two mixer inputs, until the minimum of the modulated signal

detected at the mixer output is found. Generally the sweet point is a few degrees away from the quadrature point. For some mixers, however, no sweet point is found. An interesting study of the AM sensitivity of mixers used as phase detectors can be found in (Brendel et al., 1977; Cibiel et al., 2002). The intrinsic phase noise of the Minicircuits mixer ZEM-4300 and of the Marki Microwave M80412 is shown in Figure 4.14 when both LO and RF inputs are saturated (7 dBm for the ZEM-4300 and 12 dBm for the M80412). When the power of one of the two inputs is reduced, the resulting phase noise measured at the output of the mixer increases. It should be noted though that it is not the intrinsic noise of the mixer to increase, but the sensitivity of the mixer k_d that gets smaller, leading to an higher equivalent phase noise at the output of the mixer. The increased equivalent phase noise of the Marki mixer when the RF port is driven with reduced power (-10 dBm) is also shown in the Figure.

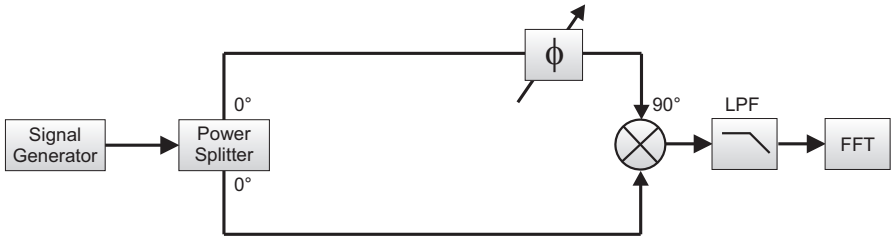

Figure 4.13: Intrinsic mixer phase noise measurement set-up.

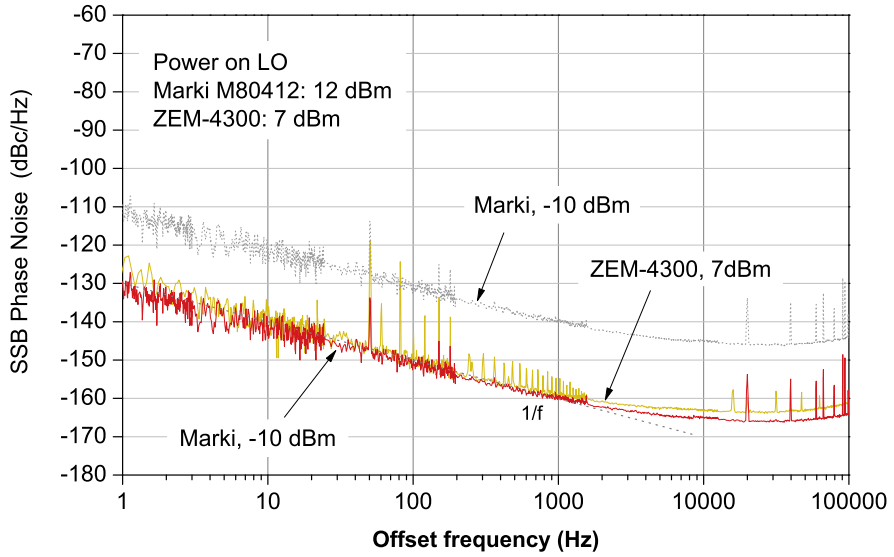

Figure 4.14: Single sideband phase noise of the ZEM-4300 mixer measured at 1.2 GHz and of the Marki M80412 measured at 8 GHz..

4.5.2 Microwave amplifier noise

In order to increase the power level of the recovered repetition rate signal to a level compatible with the typical driving levels of a microwave mixer, high gain amplification is required. Amplifiers exhibit white and flicker phase noise and for amplifiers working in the 2-8 GHz range, typical values for these noise types are usually less than −100 dBc/Hz at 1 Hz Fourier frequency for the flicker noise and −150 dBc/Hz for the white phase noise. It is worth noticing that, contrary to what is generally misleadingly reported in amplifier datasheets, the white phase noise value is not fixed but depends on the input power applied[4]. For frequencies below 2 GHz, I have chosen to use low noise amplifiers from Minicircuits Inc. ZX60-33LN which exhibit a gain of 16.6 dB at 1.5 GHz and are capable of an output power of 16.5 dBm (1 dB compression point). The phase noise of the ZX60-33LN is reported in Figure 4.16 and shows a flicker phase noise value of −125 dBc/Hz at 1 Hz and measured at 1.2 GHz. This specific amplifier was selected as it exhibited the lowest flicker phase noise among the different models that were measured. For comparison, the phase noise of two other amplifiers (JCA319 and Minicircuits ZFL-1000LN) are shown. The photodetector shot and thermal noise contribution are also shown, assuming an incident optical power of 100 nW. For frequencies between 2 and 8 GHz I have chosen to use ultra-low noise amplifiers from Hittite (model HMC-606) which exhibit a flicker noise of −120 dBc/Hz at 1 Hz as shown in Figure 4.16. The noise of these amplifiers was measured with the set-up shown in Figure 4.15.

Usually three cascaded amplifiers are needed to increase the selected harmonic of the repetition rate to a suitable driving level for the mixers. As the white phase noise of a microwave amplifier depends on the input power, the first amplifier will have the biggest contribution to the overall white phase noise. The white phase noise of the first amplifier can be calculated from the amplifier noise figure NF (1 dB for the ZX60-33LN amplifier), the typical input power level P_{in} (−40 dBm) and the thermal noise N_{R} at 290 K (−174 dBm):

$$\begin{aligned} L(f) &= N_{\mathrm{R}} - P_{\mathrm{in}} + \mathrm{NF} \\ &= -174\,\mathrm{dBm/Hz} - (-40\,\mathrm{dBm}) + 1\,\mathrm{dB} \\ &= -136\,\mathrm{dBc/Hz}. \end{aligned} \tag{4.5.1}$$

The flicker noise of the individual amplifiers is uncorrelated and therefore adds statistically. With three amplifiers, the overall flicker noise will be $20\log\sqrt{3} = 4.8$ dB larger

[4] In some amplifier datasheets and in some literature the term 'flicker noise corner frequency' is used and refers to the frequency at which the contribution from the white and flicker phase noise is the same. Unless the power level of the signal applied at the amplifier input is specified, the term is used improperly. Different power levels will in fact lead to different corner frequencies, due to the power dependence of the white phase noise.

than that of a single amplifier. For three amplifiers the flicker noise at 1 Hz offset will therefore be −120 dBc/Hz (at 1 GHz) and −115 dBc/Hz (at 8 GHz) for the ZX60-33LN and the HMC-606 amplifiers respectively.

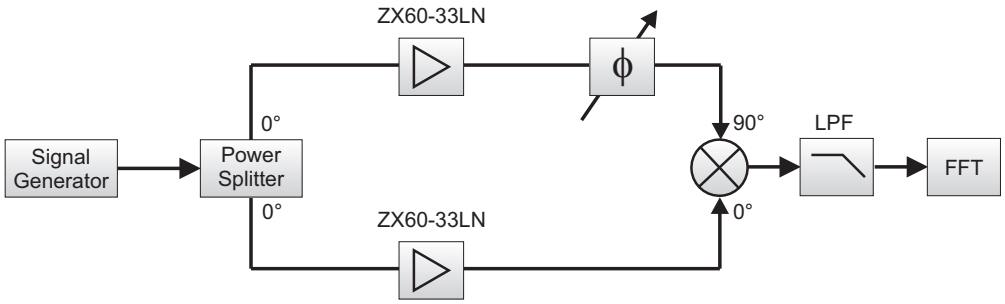

Figure 4.15: Amplifier phase noise measurement set-up.

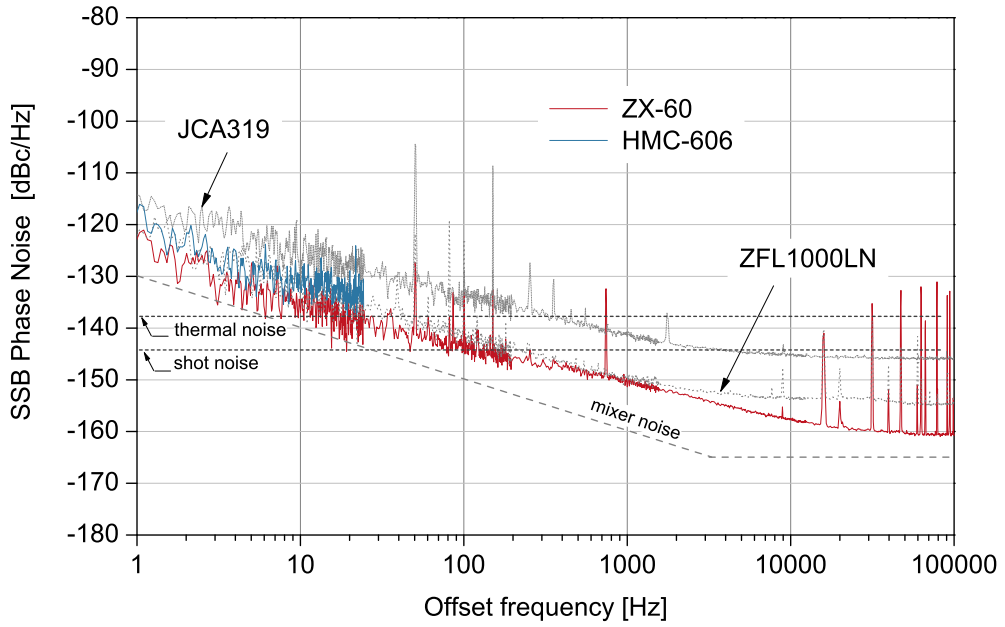

Figure 4.16: Phase noise of the microwave amplifier ZX60-33LN measured at 1.2 GHz and the HMC-606 measured at 8 GHz. The phase noise of two other amplifier models are also given, JCA319 and Minicircuits ZFL-1000LN. Although the flicker phase noise level of the two Minicircuits amplifiers is very similar, the ZX60-33LN was chosen for its higher output power level (16 dBm instead of 3 dBm), which is more compatible with typical driving levels for a microwave mixer. The higher white phase noise above a few kHz for the ZFL-1000LN is due to the lower input power used, with respect to that used with the ZX60-33LN, required to avoid saturation. The shot and thermal noise levels of the photodetector and the residual phase noise of the ZEM-4300 microwave mixer, are also shown for comparison.

4.5.3 Low frequency amplifier noise

When a mixer is used as a phase detector, the output voltage (proportional to the phase difference of the signal applied at the inputs) typically needs to be amplified by a low frequency amplifier before further processing. In order for phase information to be preserved during amplification, the noise of the amplifier must be lower than that of other components preceding it, such as the mixer or the microwave amplifiers.

The typical phase sensitivity of a mixer varies from 0.1 to 0.3 V/rad, depending on the input power levels. As a conservative choice in the analysis of the low frequency amplifier voltage noise, we can assume a value of 0.1 V/rad. The white phase noise of the amplifiers employed to increase the power level of the detected repetition rate (−136 dBc/Hz) generates a corresponding voltage noise power spectral density (PSD) at the output of the mixer of

$$\mathrm{PSD_V}(f)_{\mathrm{white}} = 10^{-\frac{136}{20}} \frac{\mathrm{rad}}{\sqrt{\mathrm{Hz}}} \cdot 0.1 \frac{\mathrm{V}}{\mathrm{rad}} = 15.8 \frac{\mathrm{nV}}{\sqrt{\mathrm{Hz}}}. \tag{4.5.2}$$

The overall flicker noise for three cascaded amplifiers is −120 dBc/Hz at 1 Hz. The PSD of the mixer output voltage due to this flicker phase noise is

$$\mathrm{PSD_V}(f)_{\mathrm{flicker}} = 10^{-\frac{120}{20}} \frac{\mathrm{rad}}{\sqrt{\mathrm{Hz}}} \cdot 0.1 \frac{\mathrm{V}}{\mathrm{rad}} = 100 \frac{\mathrm{nV}}{\sqrt{\mathrm{Hz}}}. \tag{4.5.3}$$

The typical input voltage white noise of a good operational amplifier is generally a few $\mathrm{nV}/\sqrt{\mathrm{Hz}}$ and the flicker noise at 1 Hz offset frequency is within a few tens of $\mathrm{nV}/\sqrt{\mathrm{Hz}}$. I have used the popular low noise amplifier OP27 which provides a white noise floor of 3 $\mathrm{nV}/\sqrt{\mathrm{Hz}}$ and a flicker noise of 9 $\mathrm{nV}/\sqrt{\mathrm{Hz}}$ at 1 Hz. The input current noise is 0.45 $\mathrm{pA}/\sqrt{\mathrm{Hz}}$. With the chosen mixer load resistance of 1 kΩ, the input current noise generates a voltage noise of 0.45 $\mathrm{nV}/\sqrt{\mathrm{Hz}}$. These values are well below the voltage noise levels generated at the output of the mixer by the microwave amplifiers. Therefore, the low frequency amplifier does not constitute a limiting element in the measurement set-up.

4.5.4 Phase noise budget

A summary of the contributions from each optical or electronic component to the overall white phase noise is reported below for a typical signal power of −40 dBm.

Source of noise	$\mathcal{L}(f)$ (dBc/Hz)
Photodetector shot noise	–144
Photodetector thermal noise	–137
Microwave mixer	$<$ –160
Microwave amplifier	–136

Table 4.3: White phase noise budget for a typical optical signal power of approximately 60 μW, genearing microwave signal power of -40dBm at the output of the photodector.

At low offset frequencies, the phase noise will be dominated by flicker noise. A summary of the contributions to this type of noise is reported below. The value reported is at 1 Hz offset and it is assumed that the mixer can be saturated.

Source of noise	$\mathcal{L}(f)$ (dBc/Hz)
Microwave mixer	–130
Microwave amplifiers	$\leq$–115

Table 4.4: Flicker phase noise budget.

4.5.5 Mechanical stability

When measuring stabilities at the 10^{-15} level or better for a frequency of 10^9 Hz, the mechanical stability of the electronic components and the interconnections between them becomes of paramount importance. Acoustic noise and thermal expansion affecting electronic devices, connectors and cables can in fact seriously compromise the measured short and long term phase stability of the microwave signal. The change in the signal path length corresponding to a change in phase of the microwave carrier is

$$l(t) = \frac{1}{2\pi}\frac{c_{\mu w}}{f_0}\phi(t) \tag{4.5.4}$$

where $c_{\mu w}$ is the propagation speed of the microwave signal and f_0 is the frequency of the RF carrier. For a signal path made of copper and of a typical length of 1 m at 293 K, using the thermal expansion coefficient value found in (Kaye and Laby, 1995), we have

$$\frac{\delta L}{\delta T} = 16.5\,\frac{\mu\mathrm{m}}{\mathrm{K}}. \tag{4.5.5}$$

The time delay δL caused to the RF carrier by the length change of the cable can be calculated as

$$\frac{\delta t}{\delta T} = \frac{\delta L}{\delta T}\frac{1}{c_{\mu w}} \simeq 80\,\frac{\text{fs}}{\text{K}} \tag{4.5.6}$$

and the corresponding fractional frequency instability caused by the temperature-induced time delay is

$$\frac{\delta f}{f_0}(T) = \frac{\delta t}{\delta T}\frac{1}{\tau} = 8 \times 10^{-14}\,\frac{1}{\text{K}}\,\tau^{-1} \tag{4.5.7}$$

where τ is the measurement time. The physical meaning of the above result is that, in order to achieve a frequency stability of 10^{-15} at 1 s the temperature of the signal path must be kept within $\sim$ 12 mK over that measurement time. The air temperature in the laboratory where most of the experiments described in this thesis took place (with the exception of the experiment on the Aurora network) is stabilized to within 150 mK. The required signal path temperature stability over such a short measurement time is easily achievable by just relying on the thermal mass of the cables and the microwave devices. However, if a fractional stability of 10^{-17} is expected for much longer time scales, say 100 s, keeping the temperature within 12 mK can be challenging unless the thermal mass is substantially increased. For this reason, I have rigidly mounted all the microwave devices used in the experimental set-ups on aluminium breadboards of the type generally used for optical set-ups. Thermally conductive sheets are inserted between the metal enclosures of the devices and the board to ensure a good thermal contact. The thermal contact between the cables and the breadboard is achieved through the connectors. The breadboard is also covered with a thick sheet of thermally insulating foam, mainly to stop any residual effects due to the air turbulence above it.

4.5.6 A note on frequency counters

In some of the experiments that will be described, I made use of a frequency counter to measure the transfer stability. From the theory governing the functioning of a reciprocal counter[5], the lower the frequency to be measured, the higher the resolution that can be achieved. In practice, below a certain frequency the resolution starts decreasing because of the reduced slew rate of the signal. The lower the slew rate, the less precise will be the

[5] A reciprocal frequency counter measures the period of the input signal and calculates the frequency mathematically. The period is measured as a multiple of the master clock cycle. In this way, the resolution is independent of the input frequency (not accounting for slew rate limitations). All modern frequency counters are reciprocal.

determination of the exact zero crossing point required to perform the measurement. The slew rate can be increased, to some extent, by amplification. However, strong saturation (clipping) of the amplifiers should be avoided as this could introduce unwanted phase delays. In Figure 4.17, the resolution of the Agilent 53132A counter is measured for different frequencies. In order to perform this measurement, a direct digital synthesiser locked to an H-maser was used as a test signal as its stability is better than the resolution to be measured. The best resolution is achieved for input frequencies around 50 kHz and this explains why I have chosen such frequency as the measurement beat frequency in the experiments that follow. The resolution of the counter was also measured for different slew rates of the same input signal corresponding to different amplitude levels, as shown in Figure 4.17.

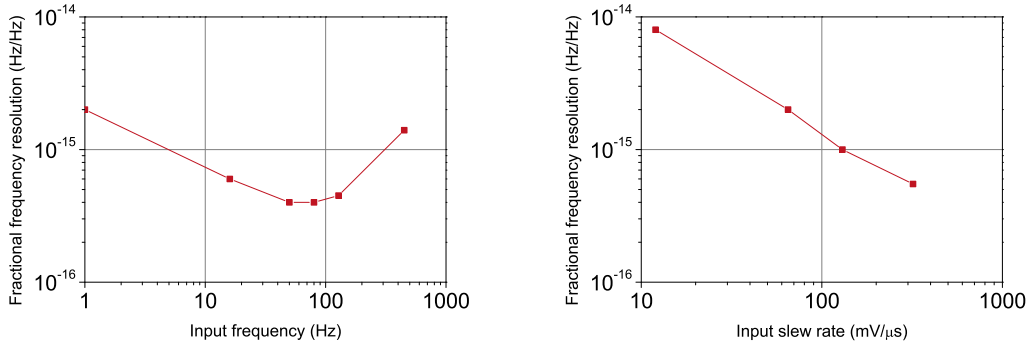

Figure 4.17: Left: Fractional frequency resolution of the counter for a gate time of 1 s as a function of the input frequency, scaled to 1.5 GHz. Right: Fractional frequency resolution of the counter for a gate time of 1 s as a function of the slew rate of the input signal, scaled to 1.5 GHz.

4.5.7 More details on the mode-locked laser module

The erbium-doped femtosecond mode-locked laser used for most of the experiments of this thesis is commercially available from Menlo Systems (model FC1500) and is part of the NPL transportable optical frequency comb routinely used for optical frequency measurements (Walton et al., 2008). For the comb transfer experiment over the Aurora network at the ORC, a newer version of the Menlo comb was used. The main difference resides in the repetition rate which is 250 MHz instead of 100 MHz. The datasheets for these mode-locked lasers do not provide exhaustive information about their internal structure, but some phone conversations I had with Menlo's technical department allowed me to confirm some basic characteristics. The mode locking is achieved by nonlinear polarization rotation (Tamura et al., 1993). Three computer-controlled piezos 'squeeze' a section of the fibre in order to change the polarization state of the laser light. A free-space polarization control section, constituted by a $\lambda/2$ and two $\lambda/4$ waveplates, complements the squeezers in order to cover all possible polarization states. The erbium-doped fibre is pumped with a single 980 nm pump diode laser. A photograph

of the laser module with the lid removed is shown in Figure 4.18, where the waveplates and other optical components are visible.

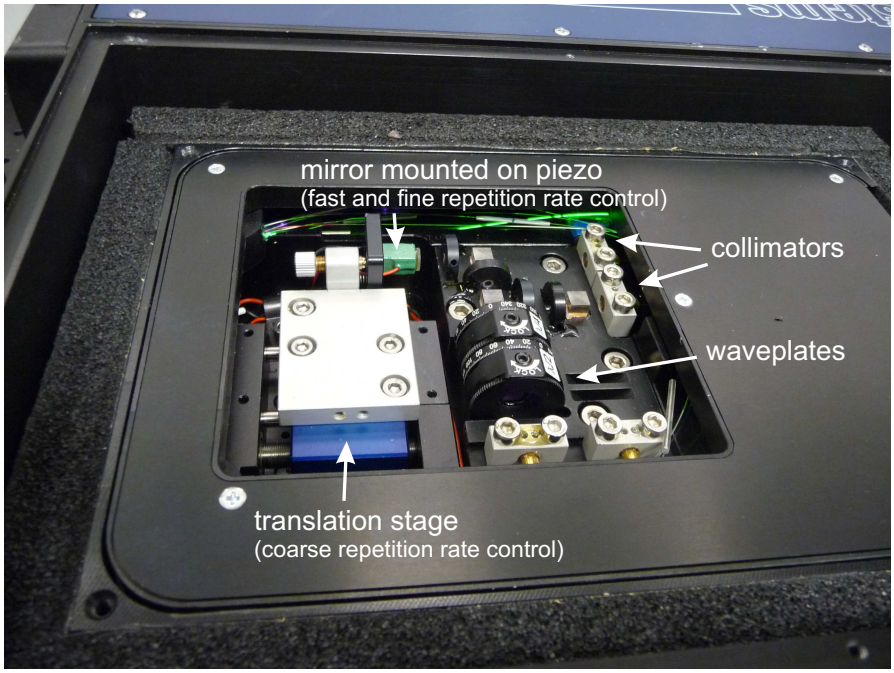

Figure 4.18: Photograph of the NPL Menlo Systems fibre laser module with the lid removed. The free space optics is required to achieve a wide tuning range of the repetition rate. The waveplates complement the computer-controlled squeezers in order to allow all possible polarization states to be obtained.

Two fibre-coupled outputs are available from the femtosecond comb system: one directly from the laser module and capable of a maximum power of 5 mW; another one after an erbium-doped optical amplifier (EDFA) that increases the optical power of the mode-locked laser up to a maximum of 150 mW. The direct output from the laser does not provide enough power to recover a suitable signal after the pulses have travelled through more than 100 km of standard and dispersion compensating fibre (∼27 dB loss), therefore the amplified output is used. The optical bandwidth of the EDFA is broadened by pulse self-phase-modulation (SPM), as the bandwidth of standard EDFAs (∼30 nm) would not be sufficient to preserve the full optical spectrum generated by the mode-locked laser. To achieve the required SPM, the pulse duration is managed within the EDFA at the design stage (Again this information comes from phone conversations as it does not appear in the system manual). The EDFA has a free space and a fibre-coupled output. The total optical power available on the amplifier is split between these two outputs in an adjustable ratio dependent on the polarization state. Also for the laser, the polarization state is changed by means of three computer-controlled squeezers pressing on a section of the fibre. The SPM process taking place within the EDFA inevitably affects the flatness of the gain of the amplifier across the optical spectrum. Different polarization states lead to fairly different optical output spectra, as shown in Figure

4.19. I could not find any evident relationship between the spectral shape and the performance of the fibre transfer experiment. Very different spectra lead to very similar transfer stability levels. This can be explained by noticing that the repetition rate is correctly extracted independent of the shape of the pulse (and hence the shape of the spectrum) as long as the pulse shape does not change with time. During the course of the experiments, many different spectra shapes have been used, as adjustments to the laser and amplifier squeezers are often required when the femtosecond comb is operated for other experiments. In particular, if the free-space output is required from the optical amplifier, the squeezer settings have to be substantially changed to divert the majority of the power to that output.

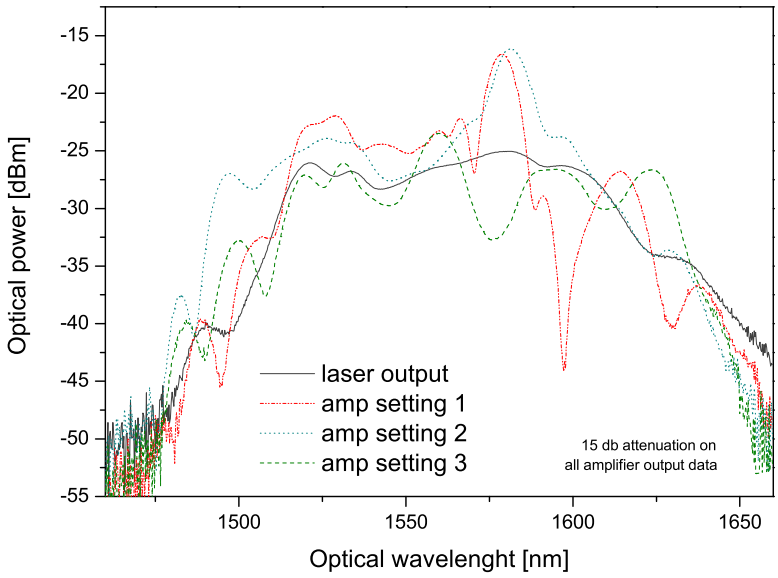

Figure 4.19: Optical amplifier output spectra for three amplifier squeezer settings of the NPL comb.

Chapter 5

Comb-based microwave transfer over 50 km of spooled fibre

In this chapter I will describe the first comb transfer experiment I performed at NPL. I this experiment I transferred a microwave frequency over 50 km of spooled fibre by propagating an optical frequency comb preserving its frequency spacing to a few parts in 10^{16} at a few tens of s when the phase noise of the fibre is cancelled. The results of this experiment have resulted in the paper "High-stability microwave frequency transfer by propagation of an optical frequency comb over 50 km of optical fiber" published in Optics Letters in 2010 (Marra et al., 2010).

5.1 Description of the experiment

The experiment performed at JILA in 2005 (Holman et al., 2005) showed that a comb could be used to transfer a microwave frequency over a many-km fibre link to a stability level that was sufficient for a selection of frequency metrology applications. However, what was measured was not the effective stability delivered at the user end but that returned to the transmitter after a round trip. The assumption of sufficient correlation between the noise accumulated by the progapating signals in the forward and backward trip was therefore required in order to be confident that the stability of the frequencies delivered to the user and after a round trip are the same. Whist this assumption could be reasonable, I believe it cannot be made *a priori* as its validity depends on enviromental conditions and an experimental validation is required. In the experiment that will be described in this chapter (and in those following it) the *actual* stability delivered at the user end was directly measured.

Unlike JILA, NPL does not yet have an optical link to other laboratories or external hubs[1]. It was therefore necessary to run this first experiment on fibre spools. Later, access to the JANET-Aurora network was granted via the University of Southampton and further comb transfer experiments could be performed on an installed network rather then spooled fibre (discussed in Chapter 6). The length of the spooled fibre link was chosen to be 50 km (one way) as a compromise between the desire to demonstrate the microwave transfer over substantially longer fibre lengths than those used in previous experiments (Holman et al., 2005, 2004) and the necessity of keeping the attenuation at reasonable levels to avoid the use of optical amplifiers as they were not available at the time.

As discussed in Chapter 3, in order for the phase noise introduced by the fibre to be measured and then cancelled, a return path for the optical signal must be created. This return path could be created by sending a portion of the optical signal back through either the same fibre or a different one lying in the same environment. In practical applications, if the fibre length exceeds 80–100 km, optical amplifiers are typically required and using the same fibre for both directions of propagation requires the optical amplifier to be bidirectional. Unfortunately, commercial optical amplifiers generally used in real fibre networks are unidirectional. Alternatively, it is possible to use a different fibre for the return path or to use the same fibre but with unidirectional optical amplifiers and circulators in break out points along the link (this latter solution will be discussed in Chapter 6). In the experiment described here I have adopted the two-fibre approach. In order for the phase noise to be suppressed efficiently there must be sufficient correlation of the noise between the forward and return paths. Here, independent fibre spools were used for the forward and return path and they were placed within the same environment. Although the correlation between the noise induced in the two paths is possibly lower than for the case of two fibres running along in the same bundle, it proved to be adequate to achieve transfer stability levels compatible with that of most of the best frequency references available today.

5.2 The experimental set-up

The set-up of the experiment can be conceptually separated into three parts:

- measurement of the phase noise introduced by the fibre;
- fibre phase noise suppression;
- measurement of the transfer frequency stability at the ‘user end’.

[1] At the time of writing NPL is putting in considerable effort in order to procure such an optical link to the external world.

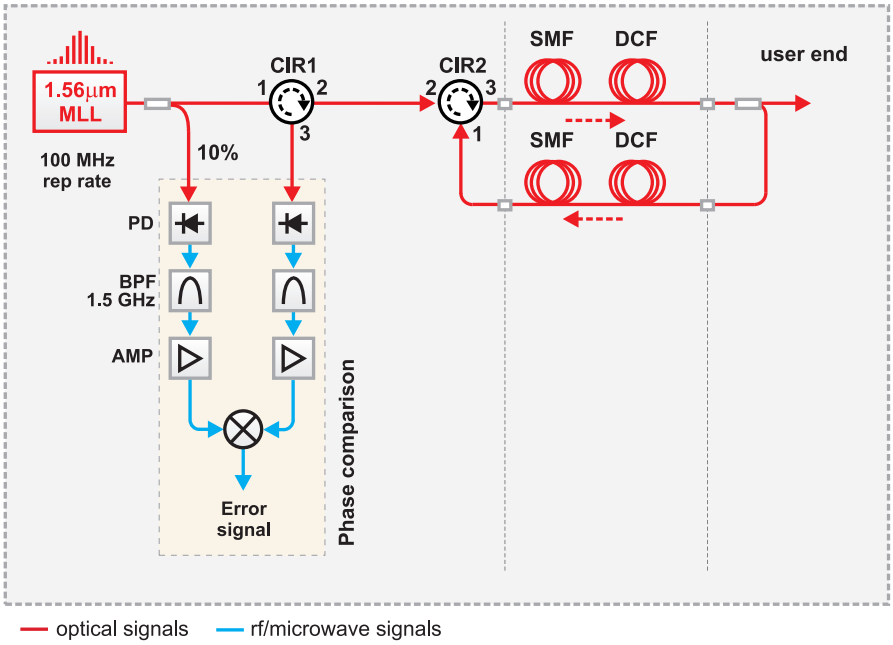

Figure 5.1: Experimental set-up for the measurement of the noise introduced by the fibre. MLL: mode-locked laser; BPF: band-pass filter; PD: photodetector; DCF: dispersion compensating fibre; SMF: single mode fibre; AMP: microwave amplifiers; CIR: circulator;

Measurement of the phase noise introduced by the fibre (figure 5.1). The measurement of the phase noise introduced by the fibre link is achieved by comparison of the repetition rate extracted by a photodetector after the optical signal has performed a round trip (100 km) with that extracted directly at the output of the laser. The phase comparison is performed using a microwave mixer with the LO and RF input signals set in phase quadrature.

Phase noise suppression (figure 5.2). Once it has been measured, the excess phase noise caused by the fibre on the repetition rate delivered at the user end is compensated by changing the optical path length of the fibre link at the transmitter end by means of a fibre stretcher.

Measurement of the transfer frequency stability (figure 5.3). Finally, the measurement of the actual transfer stability at the 'user end' is achieved by comparing the phase or the frequency of a harmonic of the repetition rate extracted with a photodetector after 50 km of fibre with that at the output of the laser.

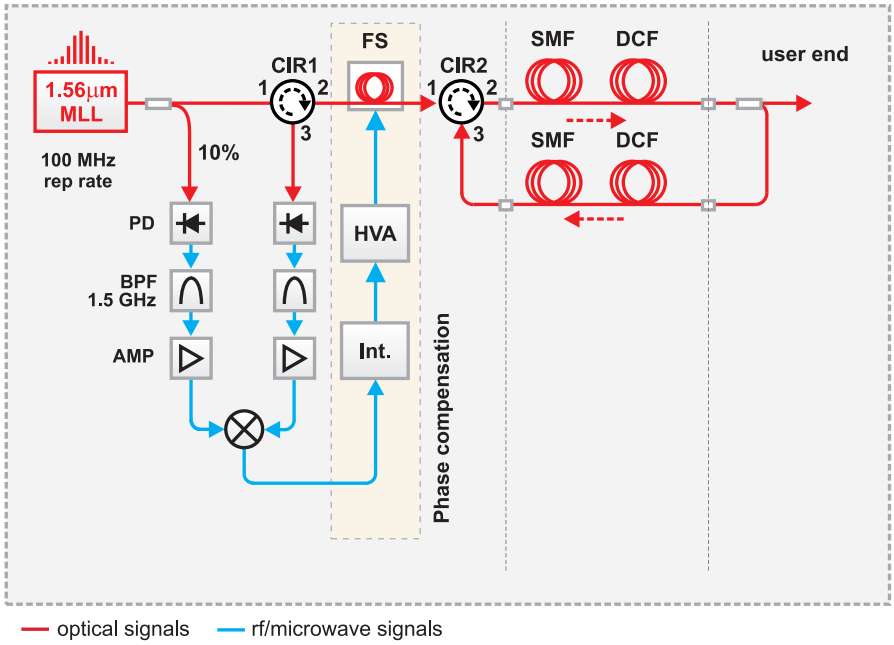

Figure 5.2: Fibre-induced phase noise compensation by means of a fibre stretcher. Int.: error signal integrator; HVA: high voltage amplifier; FS: fibre stretcher.

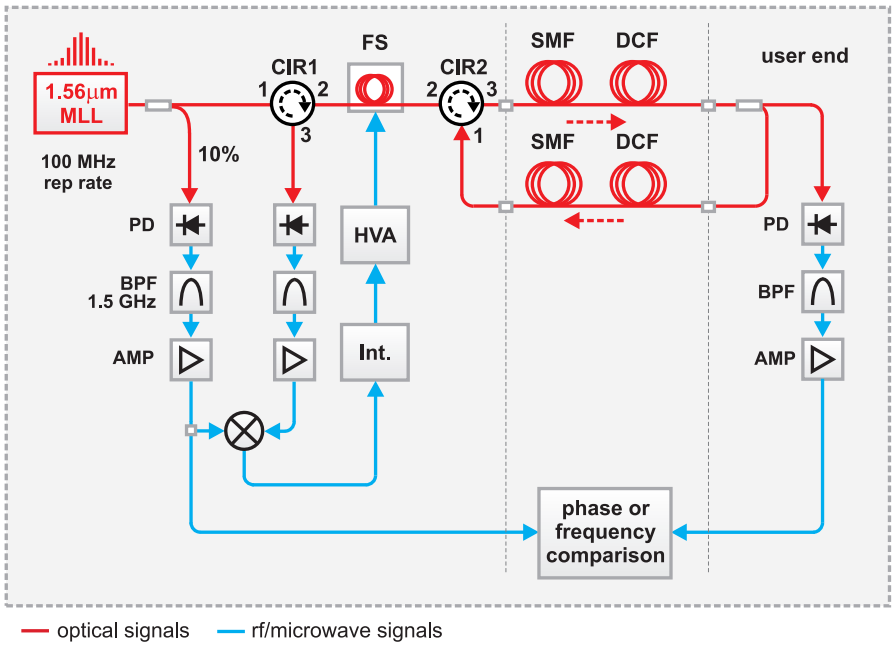

Figure 5.3: Full experimental setup for the transfer of a microwave frequency using an optical frequency comb propagating over 50 km of spooled fibre and the measurement of the transfer stability. Two separate fibres are used for each direction of propagation.

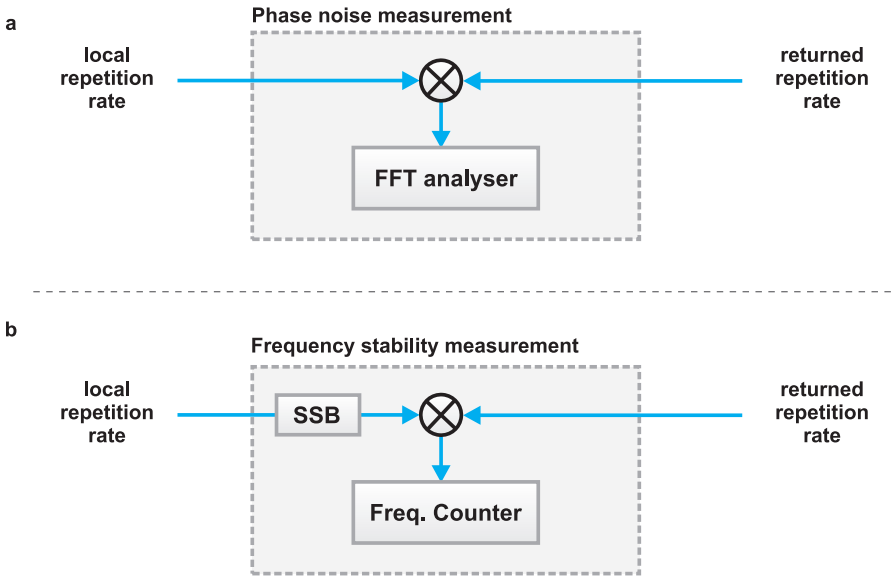

Figure 5.4: Set-up for the measurement of phase noise (a) and frequency stability (b).

5.2.1 Details of the experimental setup

The amplified output of the modelocked laser is split into two optical paths by a 90:10 splitter. After passing through an optical isolator, the high power beam is transmitted to the user end via 50 km of a standard single-mode fibre (SMF-28) and a matched dispersion compensating module (DCF). Here, another 90:10 splitter returns the larger fraction of the optical power back to the transmitter end via a second 50 km fibre and matched DCF module. The DCF module exhibits a dispersion of 864 ps/nm at 1560 nm and allows for the pulse duration to be recompressed to less than 220 ps. This pulse duration was sufficient to achieve a SNR of the extracted repetition rate, after both 50 and 100 km, in excess of 80 dB (resolution bandwidth of 1 kHz). The total round-trip optical loss was $\sim$ 30 dB of which 7 dB are due to the two DCF modules.

The low power output of the 90:10 power splitter at the transmitter end provides the low-noise reference against which the round-trip phase noise introduced by the 100 km link is measured. The laser repetition rate is detected here (local repetition rate), after a round trip (returned repetition rate), and at the user end (remote repetition rate) using battery-powered high-speed GaAs photodetectors. In each case the 15th harmonic of the repetition rate (1.5 GHz) is selected with a narrow tunable bandpass filter that suppresses the nearest harmonics by more than 30 dB. The RF power at the filter outputs is approximately –21, –40, and –47 dBm for the local, remote, and returned repetition rate signals, respectively. Two Minicircuits ZX60 low-noise cascaded microwave amplifiers (see Chapter 4 for noise details of this device) are employed for the local repetition rate and three for the remote and returned repetition rates to reach efficient driving levels for the subsequent microwave devices.

The error signal constructed by comparing the local and returned repetition rates with a microwave mixer is amplified using DC amplifiers I built using the low noise operational amplifier OP27. In addition to its low noise properties, this operational amplifier also has a very low thermal sensitivity of its offset voltages (any time-varying offset superimposed to the error signal would translate into a time-varying phase error, degrading the transfer stability). The error signal is then sent to a home built integrator and fed to a high voltage amplifier that controls the fibre stretcher. The feedback bandwidth is limited by the high voltage amplifier which was available at the time to approximately 10 Hz. With a high voltage amplifier capable of driving larger currents, the feedback bandwidth could be increased nearer to several hundreds Hz, before being limited by the 0.5 ms round-trip travel time of the pulse train. However, a larger feedback bandwidth would not help unless a better reference was used instead of an hydrogen maser, as shown in figure 5.5, as the arising self-heterodyne noise from approximately 8 Hz to 50 Hz limits the detection of the fibre noise. The limited feedback bandwidth is however sufficient to

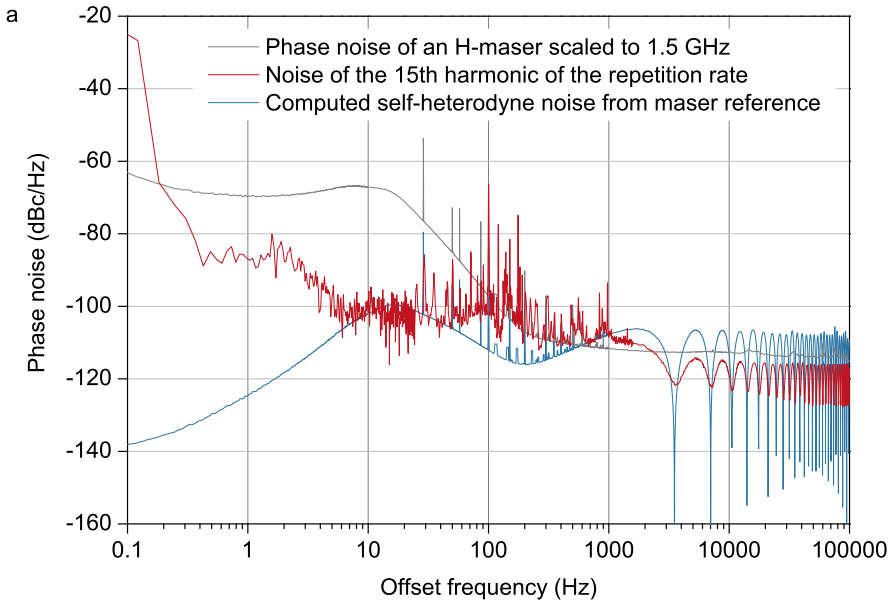

Figure 5.5: Fractional frequency stability of the 15^{th} harmonic of the repetition rate delivered at the 'user' end of the fiber with and without phase noise suppression.

efficiently compensate for the thermally-induced optical path length changes of the fibre, which become the dominant source of instability for time scales longer than 1 s. The fibre stretcher was built in-house by winding 250 turns of standard SMF-28 fiber onto a 5 cm diameter piezo tube and can be used to compensate for changes of up to 1.7 mm in the optical path length. More details on the fibre stretcher are given in Appendix A. Two optical circulators, one preceding and one following the fibre stretcher, were used to separate the pulses traveling in the forward and backward directions.

Achieving low levels of reflected power at the input of the photodetector extracting the returned repetition rate was important as no optical amplifiers were used along the optical link. The optical power available after a round trip can in fact be of a comparable level to the unwanted reflections caused by optical interconnections. Typically, a connection between two fibres using angle-polished FC/APC connectors gives reflection level between −50 dB and −40 dB. If such a connection is located in the first part of the link, the difference between the power of the returned signal after a round trip (round-trip loss 30 dB) and the reflections can be as low as 10 dB. However, the ratio between the power of the extracted harmonic of the repetition rate after a round trip to that arising from the unwanted reflections can be much bigger than the ratio between the optical powers. This is because the pulses that are reflected from any discontinuities in the fibre link (such as a connector) are broadened by the dispersion of the fibre before they reach the photodetector. In addition, the intrinsic loss of the fibre may reduce their intensity.

Therefore, the most critical connection points are those in the first sections of the link, in particular those arising from the imperfect connection of the fibre stretcher to the second circulator. For this reason I have fusion-spliced each end of the fibre stretcher to the optical circulators. I have also used angle-polished connectors everywhere else, taking extra care in making sure each connector facet was properly cleaned before connections were made.

As discussed in 4, the mechanical and thermal stabilities of the microwave electronics are also important. In particular, the dc offset that microwave mixers exhibit when operated as phase detectors, which arise from imperfect balance between the internal diodes, is temperature dependent. Thermally-induced changes in this offset can result in increased transfer instability for time scales longer than a few seconds, as the offset at the output of the mixer is indistinguishable from a genuine phase error. For this reason, I have mounted all the mixer and the other microwave components onto aluminum breadboards. The optical power splitters, isolator, and circulators, which also show non-negligible sensitivity to changes in ambient temperature, are enclosed in a box lined with thermal foam sheets as shown in figure 5.6. The measurement time was limited to 250 s (∼ 5 minutes) by the range of the fiber stretcher. The measured free running stability at 250 s is approximately $\Delta f/f = 1 \times 10^{-15}$ which corresponds to a change of the optical path length of

$$\Delta L = \frac{c}{n}\,\tau\,\frac{\Delta f}{f} = 0.6\,\mathrm{mm} \tag{5.2.1}$$

As the calculated number is an average value, the actual changes for each experimental run are larger typically by a factor of 2 to 3 such that the full 1.7 mm optical range of the fibre stretcher was found to be required most of the times.

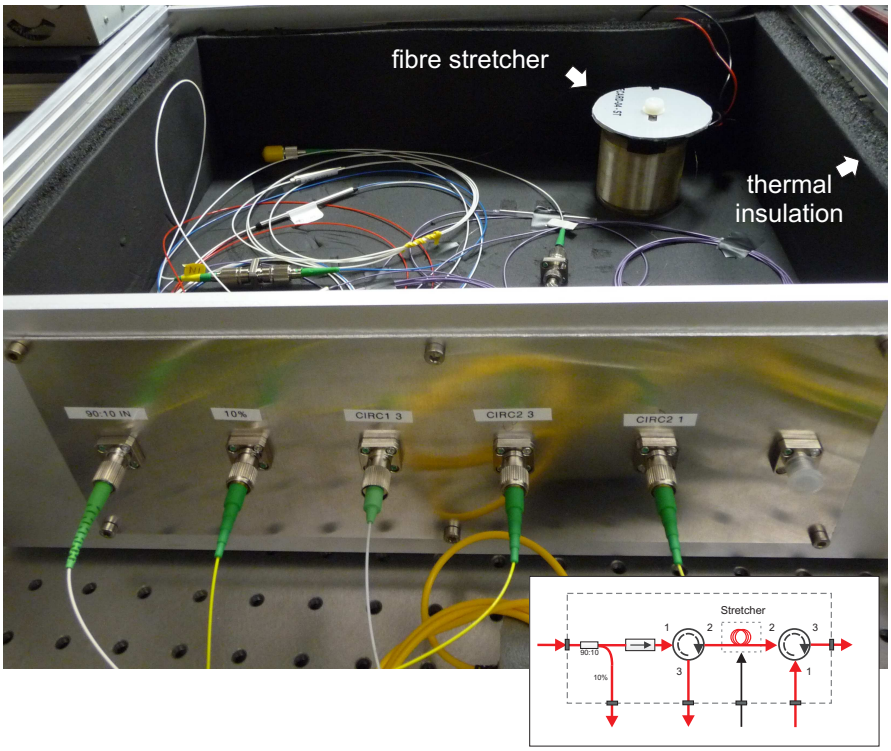

Figure 5.6: Thermal-foam lined box enclosing the fibre optics components employed in the measurement set-up in order to minimize the temperature fluctuations they are exposed to. Inset, a diagram of the optical setup in the box.

5.3 Results

The phase noise and the frequency stability of the 15^{th} harmonic of the repetition rate, measured at the user end of the fibre, are shown in figure 5.7 and figure 5.8 respectively. For the phase noise measurement, the remote repetition rate is phase compared with the local repetition rate using a microwave mixer with the inputs in quadrature and processed with an FFT analyzer up to 100 kHz offset from the carrier. With a small modification to the set-up, the fractional frequency stability can be measured. A small frequency offset of 50 kHz is added to the 15th harmonic of the local repetition rate by inserting a single sideband (SSB) modulator driven by a digital direct synthesizer (DDS), whose clock is referenced to a hydrogen maser (Fig. 5.4a). Although commercial SSB modulators can be found, I did not have one available in the laboratory and I therefore constructed the SSB modulator from discrete components. More details can be found in Appendix B.

When the phase noise cancellation loop is activated, the measured phase noise of the transferred 15^{th} harmonic of the repetition rate is suppressed by more than 10 dB for offset frequencies below a few Hz. At 1 Hz the phase noise is measured to be –101 dBc/Hz as shown in FIgure 5.7. For offset frequencies above 10 Hz no suppression is observed as the feedback bandwidth is limited to approximately 10 Hz. The self-heterodyne dips

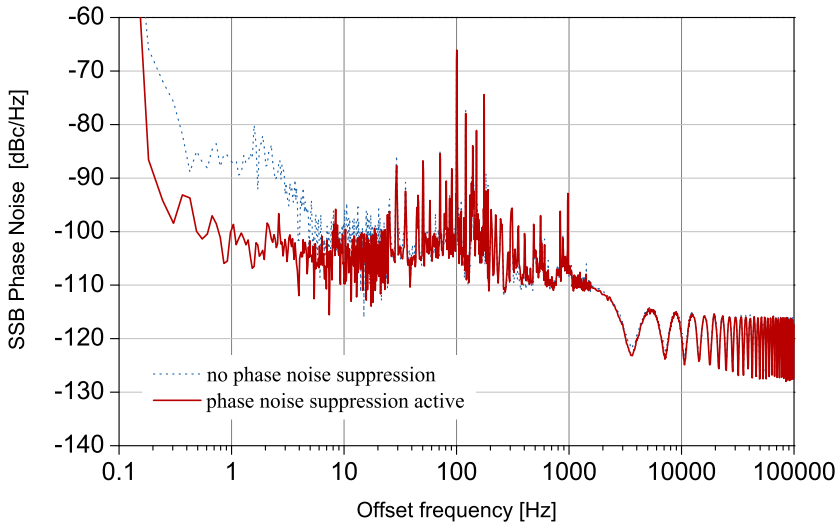

Figure 5.7: Phase noise of the extracted 15^{th} harmonic of repetition rate measured with the compensation loop on and off. The red solid and blue dashed lines are the phase noise measured at the user end with respect to the reference arm.

in the phase noise appear at multiples of 3.5 kHz, in good agreement with the values calculated in chapter 4.

The transfer stability measured at the user end is shown in figure 5.8. In this measurement each data point is obtained from independent frequency stability measurements at different counter gate times in order to overcome measurement errors due to the dead time between data samples of the frequency counter (Barnes et al., 1990). When the phase noise cancellation feedback loop is activated, the stability improves to less than 5×10^{-15} at 1 s and parts in 10^{16} for time scales longer than 10 s. For most averaging times the improvement is by a factor of 2 or more. A linear fit to the data, with the slope constrained to be –1, gives a measured transfer instability of $4.6 \times 10^{-15}\,\tau^{-1}$.

For comparison, the noise due to the electronics is also shown in the Figure 5.8. The residual instability of the SSB modulator, which needs to be sufficiently lower than that of the transferred repetition rate, was measured by driving it with a signal generator and comparing its frequency shifted output with its input and is approximately $8 \times 10^{-16}\,\tau^{-1}$.

The residual instability due to the electronics was measured by simulating the extracted 15^{th} harmonic of the repetition rate at the output of the photodetectors with a microwave signal of the same frequency and power. The simulated signal is produced by a synthesizer and is applied to both arms at the same time so that its noise is rejected. The electronics includes the filtering and amplification stages of the 15^{th} harmonic of the repetition rate, the microwave mixer used for the phase or frequency comparison,

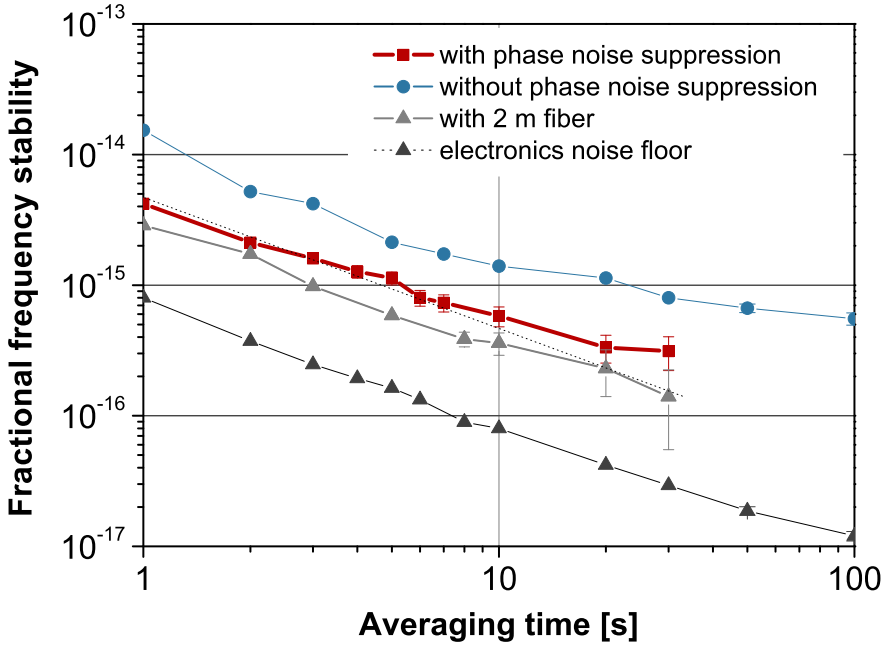

Figure 5.8: Fractional frequency stability of the 15^{th} harmonic of the repetition rate delivered at the 'user' end of the fiber with and without phase noise suppression.

the SSB modulator (for frequency stability measurements) and the low frequency amplification stages following the mixer. The results of the measurement showed the same stability as the SSB modulator, indicating that this was the dominant noise source. The measurement of the instability due to the optical part of the measurement set-up is achieved by replacing the 100 km of fibre with a short length (2 m in this case). The fibre spool losses were accounted for by inserting an appropriate optical attenuator.

The optical spectra of the comb at the input and output of the 100 km fibre link are shown in Figure 5.9.

5.4 Conclusions

I have transferred a microwave frequency over 50 km of single mode spooled optical fibre by propagation of a 90 nm-wide optical frequency comb. The transfer stability was measured to be $4.6 \times 10^{-15}\,\tau^{-1}$ showing that the technique is suitable for the dissemination of state-of-the-art frequency references as shown in Figure 5.10. In this figure, the fractional stabilities of some of the best microwave and optical references are shown together with the comb-based microwave transfer results. After a few seconds of averaging, the transfer stability becomes better or comparable to that of most of the frequency references shown. In contrast with previously published results (Holman et al., 2005), the transfer stability was measured at the user end rather than after a round trip, so that

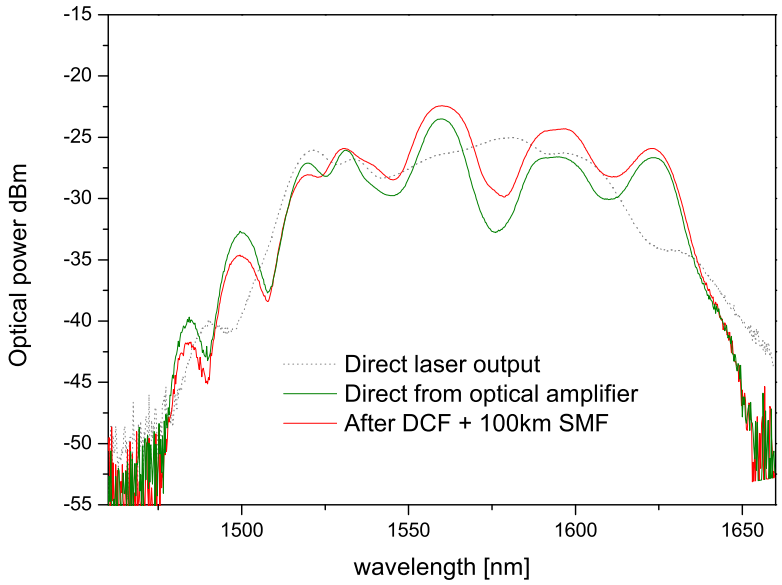

Figure 5.9: Optical spectra before and after the DCF module and 100km of SMF.

no assumptions are required to validate the results. Also, the microwave frequency was transferred using a frequency comb with an optical bandwidth of 90 nm (corresponding to over 100,000 optical modes) instead of 15 nm, and over a longer length of fibre (50 km instead of 3.45 km). However, the experiment described here was performed on spooled fibre whilst in Holman's work the fibre was part of an installed network. The knowledge acquired in preparing and running this experiment and the results achieved laid the basis for the comb-based transfer of a microwave frequency over an installed network, which is the subject of the next chapter.

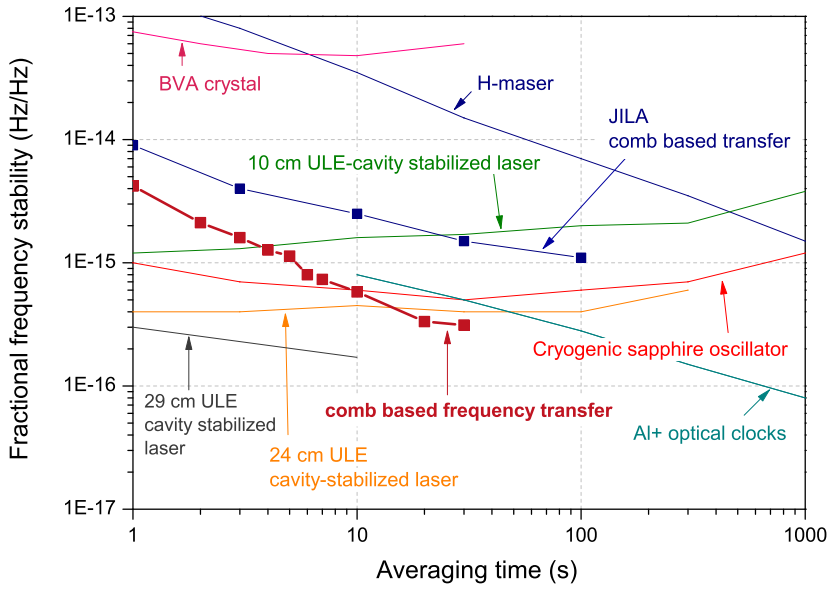

Figure 5.10: Fractional frequency stability of some state-of-the-art frequency references and the transfer stability achieved with the comb-based technique described in this chapter. 10 cm ULE cavity- stabilized laser: NPL (Webster et al., 2008); 24 cm ULE cavity-stabilized laser: NIST (Young et al., 1999); 29 cm ULE cavity-stabilized laser: NIST-Univ. of Colorado (Jiang et al., 2011); Cryogenic sapphire oscillator: NMIJ-UWA (Watabe et al., 2006); Al^{+}/Hg^{+} optical clock: NIST (Chou et al., 2010); H-maser: from datasheet; BVA crystal: from datasheet;

Chapter 6

Microwave transfer over an installed fibre

In the previous chapter I have described the microwave transfer experiment over 50 km of spooled fibre in the laboratory and demonstrated that the stability of the repetition rate can be preserved over this length of fibre at a level interesting for the frequency metrology community. The natural evolution of this experiment is to demonstrate whether the same levels of stability can be achieved over an installed fibre, where higher noise is to be expected compared with a controlled laboratory environment. This chapter describes the microwave transfer experiment I performed over the JANET-Aurora network, a research fibre network linking five universities in England. At the end of this chapter I will show that it is possible to achieve similar levels of transfer stability on installed fibre to those obtained on fibre spools and that the transfer technique is compatible with the best frequency references available at the time of writing this thesis.

This experiment resulted in the paper "High-resolution microwave frequency transfer over an 86-km-long optical fiber network using a mode-locked laser" published in Optics Letters in 2011 (Marra et al., 2011).

6.1 From spooled to installed fibre: what to expect

As previously addressed, the phase noise arising during the propagation of a frequency comb (and more generally any optical signal) through a long fibre is primarily due to the mechanical and thermal stimulations of the fibre itself by the environment. The magnitude and the spectral distribution of the induced phase noise can vary subtantially between an installed fibre and spooled fibre. In the following sections these differences are discussed.

6.1.1 Acoustically induced noise in an installed fibre

Because of the very high installation costs of dedicated ducting, fibres are often installed in existing ducts, such as those for electrical cable or gas, or on occasion even sewer lines (Gokhale, 2006). The fibre cables are installed in polyethylene microducts of different diameter and installed in the main duct, which is made of concrete or plastic materials. Such ducts, when not in a metropolitan area, often run along motorways and railway tracks. We should therefore expect excess noise on the trasmitted optical signal with a spectral distribution somewhat related to that of urban, road and rail traffic noise.

Urban noise The background noise of metropolitan areas is generally dominated by that generated by road traffic, the spectrum of which is discussed next. Apart from road traffic noise, other sources can be underground railways and central heating plants beneath buildings with their characteristic "humming" noise.

Road traffic noise Research in this field has shown that road traffic exhibits a prominent peak betwen 700 and 1300 Hz(Sandberg, 2005; Gjestland, 2008) (which unfortunately for us all falls very close to the region of maximum sensitivity of human hearing!). Another peak around 50-100 Hz is also often found.

Rail traffic noise Noise produced by trains is predominantly in the 500 to 4000 Hz region of the frequency spectrum (Thompson, 2008). This noise originates primarily from traction (with the highest level of noise generated by diesel engines), fixed installations such as level crossings and shunting noise such as the impacts between train cars during acceleration or deceleration.

The actual spectral distribution of the noise reaching the fibre can be substantially modified during the propagation of the acoustic waves through the ground, with increasing attenuation of the higher frequencies as the distance between the source of noise and the fibre increases. As far as the noise intensity is concerned, the overall noise intensity will depend on the density of human activity near the fibre. Scientists studying urban noise have found that the noise intensity is proportional to the logarithm (base 10) of the population density (Gjestland, 2008). We should therefore expect that fibre networks in densely populated countries such as Germany, England and Italy are likely to be noisier than those installed in France, Spain and Norway (an informative population density map is shown in (Gjestland, 2008)).

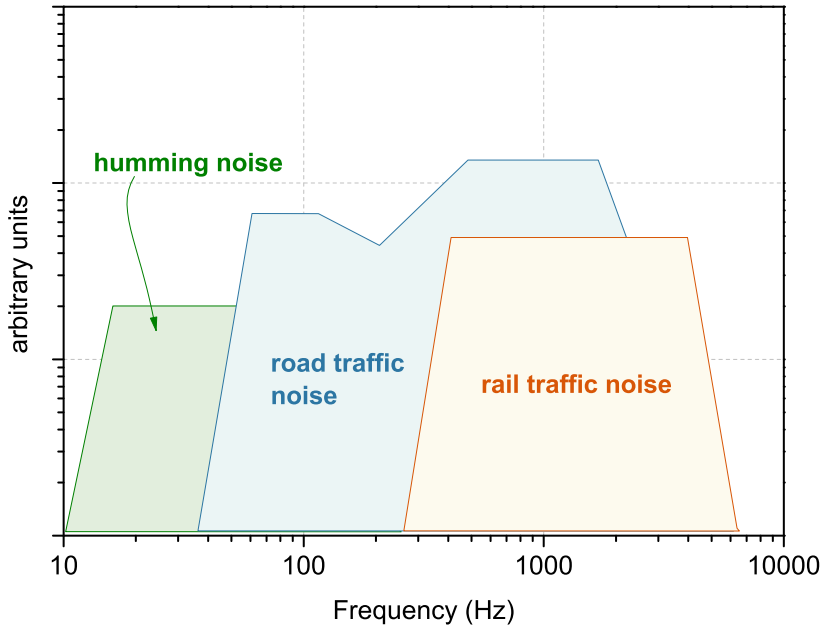

Figure 6.1: Typical frequency spectra of acoustic noise sources that are likely to induce phase noise into an underground fibre.

6.1.2 Noise induced by temperature fluctuations

Over time scales longer than a few seconds, temperature variations of the fibre will usually be the the main contribution to optical path length changes. The level of temperature fluctuations that the fibre will experience strongly depend on the burial depth. Experiments performed at the Goldstone Deep Space Communications Center (GDSCC) in California's Mojave Desert (Calhoun et al., 2002) have shown that even 1 m of ground is sufficient to suppress the surface temperature fluctuations by a very substantial amount. Table 6.1 shows the measured attenuation for different burial depths:

Burial depth (m)	Attenuation (dB)
0.9	21
1.2	30
1.5	45

Table 6.1: Attenuation of temperature fluctuations as a function of burial depth.

Using these values we can estimate the optical path length changes that we should expect on fibres installed in England. Temperature excursions during the day-night cycle in England are on average less than 15 °C. If the fibre is buried 1.5 m underground, it will experience a temperature change of no more than $4.7{\times}10^{-4}$ °C. This corresponds to a maximum change in the optical path length of only $\mathrm{d}L/L = 7$ ppm/°C $\times$ 4.7 $\times 10^{-4}$ ∘C = 3.3 $\times 10^{-9}$ over long time scales. Over shorter time scales, between a fraction of a second

and a few hundreds of seconds, the attenuation of the surface temperature variations will be much higher than that measured at the GDSCC because of the inherent inertia to changes provided by the thermal mass of the ground between the surface and the fibre. The large thermal mass effectively acts as a low pass filter for fast temperature changes.

Somewhat unexpectedly, an important role is played by the short lengths where the fibre resurfaces at breakpoints, such as manholes or above ground hubs, and the beginning and the end of the fibre link. Although these sections are generally only a small fraction of the total length of the fibre, they are exposed to surface temperature changes with much reduced attenuation. In the GDSCC experiment, the main contribution to the phase instability was indeed due to the temperature variations of the short fibre (10 m) contained in four vaults along a 16 km optical link. Although the daily temperature variations in Europe are much smaller than those in the Mojave desert (which can be as large as 50 degrees Celsius), we should still expect some contribution to the phase instability of the transmitted signal if surface breakout points are present along the fibre link. By means of an example, let us assume that a 100 km-long fibre link in England has a breakout point in the middle. The hub is located in a rural area and exposed to day/night temperature variations of 15 °C. Approximately 20 meters of fibre are used to make the connection in the hub. The optical path length change the 20 m of fibre will experience is dL = 7 ppm/°C × 15 °C × 20 m = 2.1 mm. If the fibre is buried 1.5 m underground for the remaining length, the optical path length change for this section is $dL = 7$ ppm/°C $\times 4.7 \times 10^{-4}$°C $\times 10^5$ m = 0.32 mm. Just 20 m of fibre in the breakout hub will therefore play a far more important role than the fibre buried underground.

The above considerations might not be applicable in the case of a fibre installed in a gas duct, where the temperature changes might derive from the gas flow rather than the surface temperature.

6.1.3 Spatial distribution of the environmental stimulations

The stimulations that an installed fibre receives are not only likely to be higher than those in a controlled laboratory environment, but also the effect on the phase of the optical signal travelling along the fibre will be different of the fibre is laid out straight rather than being wound on a spool.

It is interesting to note in fact that in a spooled fibre a stimulus is applied with some periodicity over the length of fibre. Using a simplified model[1], let us analyse the case of a plane acoustic wave coming from a radial direction to the fibre spool (Figure 6.2).

[1]The model does not take into account many second order effects such as the way the acoustic pressure propagates through the fibre and between the several layers of the spool, or acoustic reflections by the surfaces surrounding the spool. However, the simplified model allows the spectral region of the phase changes to be estimated

The pressure of the sound wave will "squeeze" the fibre and thus change its optical path length. The phase of the optical signal inside the fibre will be affected multiple times by a single stimulus if the latter is applied for times longer than the pulse train transit time over a turn of the spool. In contrast, an optical signal travelling through an installed fibre will generally meet the environmental disturbances only once during its transit time. This difference should in principle reflect itself in the spectral distribution of the induced phase noise.

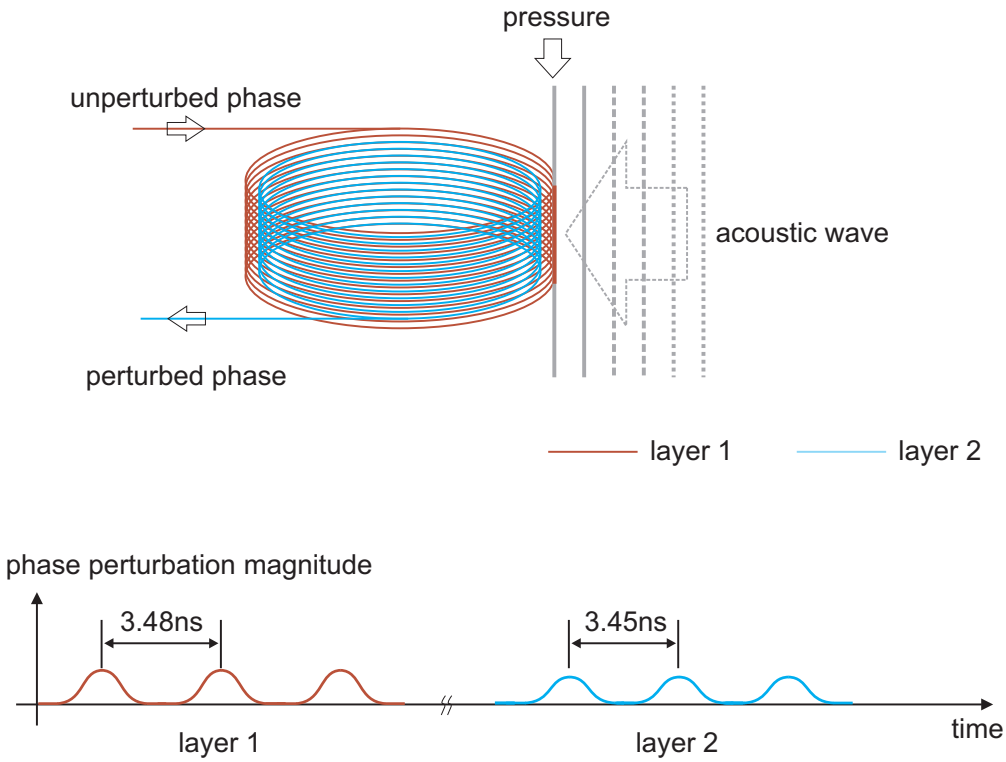

Figure 6.2: Illustration of the periodic fast phase perturbations on an optical signal caused by an acoustic wave front reaching a fibre spool from a radial direction. Only two layers are represented for simplicity. At each turn the travelling optical pulse train experiences a phase perturbation caused by the pressure of the acoustic wave on the fibre. This situation is very different from the case of an installed fibre where generally perturbations are met only once along the length of the fibre.

Fast acoustically-induced phase fluctuations We can estimate what region of the spectral distribution of the phase changes we should expect to be affected by the described effect, by calculating the repetition rate of the phase changes, starting from the first layer of windings. The outer diameter of a standard single mode fibre SMF-28 is 250 μm and the typical height of a fibre spool is 95 mm. We can therefore expect approximately 95 / 0.250 = 380 turns in the first layer, assuming tight and precise winding. The radius of the first layer is approximately 110 mm, leading to a total length for 380 turns of approximately $2\pi \times 0.11$ m $\times$ 380 = 262 m. The transit time of the

optical signal over this length of fibre is 1.3 μs. In order to have more than one acoustic wavefront within the transit time, the acoustic frequency should be higher than 1 / 1.3 μs = 770 kHz. This frequency is more than two orders of magnitude higher than typical acoustic frequencies. We can therefore consider the acoustic wavefront as a constant localized radial pressure on the fibre during the transit time of the optical signal through the first layer of the spool. The phase of the optical signal will be affected at every turn, as depicted in Figure 6.2, with a maximum deviation every 3.48 ns, corresponding to a repetition frequency of approximately 1 / 3.48 ns $\approx$ 290 MHz. The phase of the optical signal will also be affected at all the subharmonics of this frequency, with the lowest being at 290 MHz / 380 turns = 760 kHz. The wave front will of course propagate through inner layers as well, with increasing attenuation, whilst the repetition frequency of the phase change will increase up to 425 MHz for the innermost layer (75 mm diameter), with the lowest harmonic for this layer being at 1.1 MHz. Phase noise between 0.76 and 400 MHz from the carrier should therefore be expected. However, in the actual comb transfer experiment, the need for the pulse train to be returned to the transmitter end for phase noise cancellation causes the laser noise to be detected at the user end via the self-heterodyne effect. The noise starts to appear for frequencies close to the inverse of the pulse train transit time through the fibre link. For fibre lengths longer than 270 m the self-heterodyne frequency is already lower than the minimum noise frequency due to the fast acoustically-induced phase flutuations calculated above (760 kHz) and the laser noise is likely to mask the fibre noise. Also, the phase noise at offset frequencies above 100 kHz is often of little or no interest in frequency metrology applications and I will therefore not investigate this high frequency noise further and instead will concentrate on the noise that the phase perturbations can cause at lower offset frequencies via the build up effect discussed next.

Multiplicative effect of the acoustic perturbations Above we have analysed how an acoustic wave front can create a periodic modulation of the phase of the optical signal in a fibre spool due to the periodic structure of the fibre spool geometry. We have seen that these phase fluctuations are expected to be faster than several hundred kHz and therefore of negligible interest for the aim of this experiment. However, I now will analyse how these small fast phase changes can lead to macroscopic phase changes at the frequency of the acoustic wave stimulating the fibre. Each fast phase change can be seen as a small "kick" to the phase of the optical signal. As described earlier, the repetition rate of these small kicks is much higher than acoustic frequencies. Acoustic waves of frequency lower than half of the transit time of the optical signal through the spool will impose phase kicks all of the same sign. For a standard 25.2 km spool the transit time is 126μs. For acoustic frequencies up to approximately 1 / (2 $\times$126μs) = 4 kHz, the phase kicks will have all the same sign and will show as an accumulated phase shift at the end of the fibre. The outcome of this analysis is that we should expect a fibre spool to be much more sensitive to acoustic noise than an installed fibre. An experiment to compare

the spectral distribution of the phase changes a spooled fibre and an unwound fibre are compared when subject to exactly the same acoustic perturbations is far from being easy to implement. However, as I will show later, when measuring the phase noise of the installed fibre, I did find that the acoustic noise was higher for the spooled case, despite the fact that the spools were kept in a quiet laboratory environment. The multiplication effect described above could be one possible explanation for the higher noise measured.

6.2 The test bed: JANET-Aurora fibre network

JANET-Aurora is a dark fibre[2] network that links four universities in England over a total length of more than 500 km. The five linked univeristies are: the University of Southampton, the University College London (UCL), the University of Essex, the University of Cambridge University. A map of the Aurora network is shown in Figure 6.3. The comb transfer experiment described in this chapter was carried out on the Southampton-UCL branch, from the ORC to the first hub 43 km away in Crawley Court. Two fibres are available in this section and they were joined together in the hub in order to form a loop back to the ORC. This arrangement is required to allow the stability of the signal delivered at the 'user' end to be measured. The repetition rate detected after 86 km is compared with that detected before the optical comb is propagated through the fibre.

Optical time-domain reflectometry measurements performed by the ORC staff upon connection to the JANET-Aurora network show that 33 fibres with different loss coefficients per unit length are used in the ORC-Crawley Court-ORC link. Many tens of metres of fibre run within the building, thus being subject to large temperature variations and acoustic noise.

6.3 Laboratory conditions

The laboratory in which the experiment is performed is located on a raised ground floor of the ORC. A large blacked-out window is present on one of the walls. The ambient temperature is regulated with a wall mounted air conditioning unit which operates in an on-off mode, with a typical cycle of several minutes. When the air conditioning is turned off, the temperature was measured to rise up to 27 degrees Celsius because of the heat generated by the electronic and optical equipment. Although a lower temperature would have been more comfortable, I judged a slow monotonic temperature rise to be preferable to a cyclic change over several minutes, so the experiment was run with the air conditioning turned off.

[2]A dark fibre is a deployed fibre that is unused by the provider. No light is sent through the fibre, therefore the fibre is 'dark'.

Figure 6.3: The JANET-Aurora network is a dark fibre network linking four universities in the south of England over a total span of over 800 km. Two fibres are available on each the span between universities. A number of hubs are present between major spans. The comb transfer experiment described in this chapter was carried out on the first span from the University of Southampton to a hub in Crawley Court located 43 km away. The two fibres were joined together at the hub in order to create a loop where both the transmitter and receiver (user) ends are in the laboratory at the University of Southampton.

The vibration levels of the laboratory have not been measured, but they can be expected to be higher than those experienced at a laboratory such as NPL that was specifically designed to provide a low vibration environment. The main sources of noise are the fans of the electronic equipment, the activity in neighbouring laboratories, the voices and foot steps in the outside corridor and to some extent the road that runs a few tens of metres from the ORC laboratory, although its traffic is farly low. Heavy building work took place for a few days during the course of the experiment immediately outside the laboratory (a car park was built), but final results were taken outside working hours and at weekends. The whole optical setup, excluding the fibre comb, and part of the electronics were constructed on an aluminium breadboard resting on a non-optical bench.

6.4 Experimental setup differences from the spooled case

The setup for the comb transfer experiment on the JANET-Aurora network requires two major modifications to reflect the different arrangement of the fibre link:

Higher attenuation The fibre link is now 86 km long rather than 50 km. In the case of an SMF-28 type of fibre, the extra length would impose an additional 7.2 dB attenuation of the delivered comb, giving a total transmitter-receiver loss of approximately 17 dB. However, many sections of the link exhibit higher losses and additional loss is present due to imperfect splice joints, increasing the total loss of the 86-km link to 26 dB, 9 dB higher than the loss expected for an uninterrupted length of SMF-28. At the transmitter end, the returned optical comb is therefore attenuated by 52 dB (therefore 26 dB higher than for the 100 km spooled loop), becoming comparable to the level of reflections expected from the connectors used to interface the various optical components. One or more optical amplifiers are therefore required. The consequence of the insertion of the optical amplifier is the optical bandwidth limitation that will be imposed onto the transmitted frequency comb. Whilst in the experiment on fibre spools the optical bandwidth was as high as 100 nm, the insertion of an optical amplifier limits it to approximately 30 nm in this experiment.

Return path via the same fibre Only two fibres are available between the ORC and the hub at Crawley Court. The test of the true stability delivered at the 'user' end of the fibre link means that the 'user' end must be located in the laboratory. Therefore, the optical comb needs to be returned via the same fibre rather than a separate one as in the case of the spooled fibre experiment.

As discussed earlier, at least one optical amplifier is required. However, commercial optical amplifiers are unidirectional whilst bidirectional amplification is required for this experiment as the comb is propagated in the same fibre in both directions. Two solutions are possible to overcome this problem. One solution is to separate the forward from the backward travelling pulse train using two circulators and to use a unidirectional amplifier on each arm. This arrangement can be adopted at any breakout point along the link, in this case the hub in Crawley Court. The commercial optical amplifiers used in this experiment allow for the amplification parameters (such as gain, output power,etc..) to be remotely controlled via the Internet Protocol, making this solution a viable one. The other solution is to insert the optical amplifier at the 'user' end, again by means of an optical circulator. These two possible arrangements are illustrated in Figure 6.4.

The latter is a much more preferable solution as it requires only one amplifier and one circulator, with the additional advantage that the amplifier is located in the laboratory and therefore under full operational control. Also, temperature variations in the laboratory are likely to be smaller and more controllable than those in a remote hub. The temperature stability of the circulator, amplifier and the fibres interconnecting them are particularly important as the pulse train propagates through them unidirectionally. Imbalances between the temperature-induced phase fluctuations in this part of the set-up

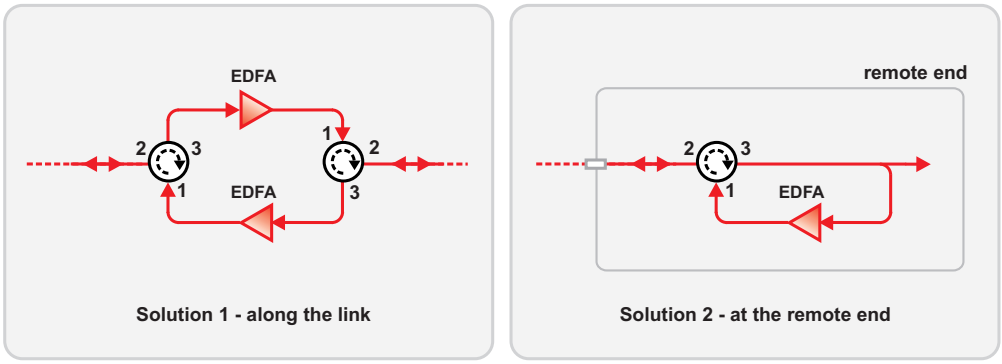

Figure 6.4: Possible solutions to overcome the unidirectionality problem with optical amplifiers. The solution on the left of the figure can be adopted anywhere along the link, but requires two amplifiers and two circulators and remote access to the parameters of the amplifiers. The solution on the right of the figure is applicable at the user end of the link and it preferrable as it employs only one amplifier and one circulator and no remote access is required.

could greatly contribute to the instability of the repetition rate as they are not common-mode fluctuations unlike the case where both pulse trains propagate bidirectionally in the same fibre. This situation is illustrated in Figure 6.5.

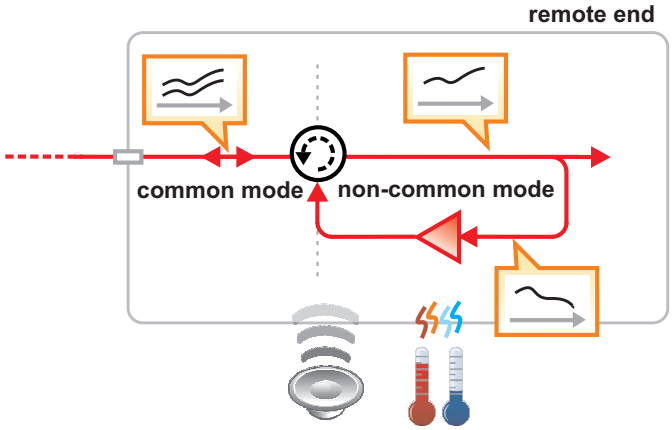

Figure 6.5: Any section of the optical setup where the pulse train propagates in different fibres for each direction can be a source of instability. Noise deriving from the acoustic and thermal effects acting on the non-common-mode part of the optical path cannot be cancelled.

Thermally controlled fibre spool In the previous experiment on fibre spools the measurement time was limited to a few minutes due to the small optical path length compensation range of the fibre stretcher. For this experiment on installed fibre, two rather than one fibre stretchers are used and a thermally controlled 300 m-long fibre spool is inserted to extend the range to 50 mm. The fibre stretchers are controlled with a relatively fast feedback loop (10 Hz) whilst a slow (many seconds) feedback loop

controls the thermally controlled spool. This arrangement allows for the large optical path length changes, due to long term temperature fluctuations of the fibre link, to be compensated.

Other differences, not imposed by the fibre network arrangement, characterize the new setup compared to the setup used in the fibre spool experiment.

Higher harmonic repetition rate detection Whilst the phase noise of a harmonic of the detected repetition rate increases with harmonic number, the amplitude noise is frequency independent. Amplitude noise becomes a problem when the repetition rate signal is applied to a mixer operated as a phase detector, as this has still a non-zero amplitude sensitivity (see Chapter 4). At the output of the mixer, voltage changes due to amplitude and phase changes are indistinguishable. Choosing higher harmonics of the repetition rate reduces the contribution to the output voltage of the mixer due to amplitude noise with respect to that due to phase noise. In the experiment, three fast photodetectors would be needed for the detection of the 'local', 'returned' and 'remote' repetition rate. However, only two fast photodetectors were available at the time of this experiment, so I could only partially exploit the benefits of a reduced AM sensitivity by using higher harmonics of the repetition rate. The higher detection frequency also enables lower microwave noise floors. The phase noise of microwave amplifiers working at 1.5 and 8 GHz can be very similar, thus making the phase noise relative to the carrier (dBc) a factor of $20\log(8/1.5) = 14.5$ dB smaller for the 8 GHz case.

Higher repetition rate laser The fibre laser used in the network experiment is from the same manufacturer (Menlo Systems, Germany) as that used at NPL for the spooled fibre experiment, but it is a newer version. The main difference of this laser is its repetition rate which is 250 MHz instead of 100 MHz. For the purpose of this experiment, this difference does not have major consequences, apart from a smaller number of optical modes being transmitted per unit optical bandwidth. Another difference is the larger optical spectrum that is covered after the output of the laser is amplified by the nonlinear amplifier. The repetition rate and the carrier offset frequency are phase locked to synthesizers referenced to a GPS-disciplined oven-stabilized crystal oscillator.

6.5 Experimental setup

The experimental setup is shown in Figure 6.6. The mode-locked laser generates a pulse train with repetition frequency of 250 MHz and a pulse duration of approximately 150 fs. The fibre laser rests on a bench on the other side of the lab with respect to the experimental optical setup. The laser and the setup are linked using a 6 m-long SMF-28 optical fibre running along the walls. From the dispersion value of SMF-28 fibre at 1560

nm (17 ps/km · nm) it is estimated that the pulses are as long as 40 ps before they enter the experimental set-up.

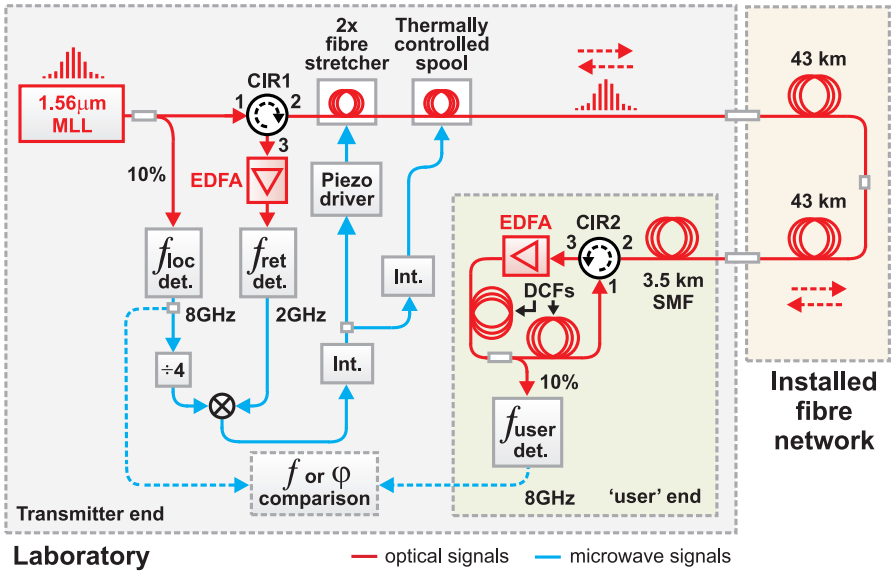

Figure 6.6: Experimental set-up for the transfer of a microwave frequency using an optical frequency comb over the JANET-Aurora network. Each repetition rate detection module consists of a photodiode followed by a tunable band pass filter and cascaded RF amplifiers. The block marked as f or ϕ comparison is either a phase detector followed by a fast Fourier transform analyser or a single sideband modulator-based frequency counting system. MLL: mode-locked laser; CIR: optical circulator; SMF: single mode fibre; Int.: integrator; Det.: detection stage; DCF: dispersion compensating fibre; EDFA: erbium-doped fibre amplifier.

6.5.1 Principle of operation

The pulse train is initially split in two by a 90:10 optical power splitter. The lower power output is used to generate the reference repetition rate against which both the 'remote' and 'returned' repetition rates will be measured. The high power output is connected to port 1 of an optical circulator, whose port 2 is connected to two cascaded fibre stretchers and a thermally controlled spool. The fibre stretchers and the thermally controlled spool are the actuators for the phase noise cancellation loop described later. The output of the thermally controlled spool is connected to the JANET-Aurora fibre, extending to Crawley Court and then back to the laboratory. Here, a 3.5 km-long fibre spool is added to the network in order to improve the matching of the dispersion compensating fibre modules used to achieve suitably short pulses for both the 'user' and the 'returned' pulse trains. The circulator CIR2 is used to allow for an unidirectional amplifier to be inserted at this point of the fibre link. The incoming dispersed pulse train available at port 3 of the circulator is amplified and then recompressed using a dispersion compensating module. The output for the 'user' is obtained at the output of

the module with a 90:10 splitter. The high power output is sent back to the network via another dispersion compensating module, used to pre-compensate for the dispersion encountered during the return trip, connected to port 1 of the CIR2. At the transmitter end, the returned pulse train is extracted at port 3 of CIR1 and amplified.

A harmonic of the repetition rate is detected from the pulse train directly from the laser, at the 'user' end and after a round trip through the optical link. I will refer to them as 'local', 'remote' and 'returned' repetition rate as for the previous experiment. At the transmitter end the local and returned repetition rates are phase compared to construct an error signal which, after integration, provides a control signal for the actuators (fibre stretchers and thermally controlled spools). The test of the actual phase and frequency stability delivered at the user end is achieved by comparing the remote repetition rate with the local repetition rate directly available at the output of the laser.

Below follow more detail each subsection of the measurement set-up:

Fibre stretchers Unlike the previous experiment on spooled fibre, where a single home-made fibre stretcher was employed, in the experiment on the installed network I used two commercial fibre stretchers. These are manufactured by Opti-Phase and consist of a piezo tube of 6 cm diameter on which 60 m of fibre are tighly wound over two layers. The total optical path length change achieved via these two stretchers is 12.8 mm.

Thermally controlled spool I constructed this extended optical path length range actuator by manually winding 300 m of fibre onto an elliptical tube made out of copper sheets. An illustration of the design of the thermally controlled spool is shown in Figure 6.7. Two Peltier heaters/coolers apply the temperature changes to the fibre via the copper sheet bracing part of the spool, as illustrated in Figure 6.7. Soft thermally conductive sheets are used to improve the thermal contact with the fibre and the copper sheets. I have chosen the elliptical shape as the majority of the fibre is exposed with a large curvature radius, making the contact between the heating/cooling copper sheet easier and more efficient. Large heat sinks are used to dissipate the heat when the Peltier actuators are heating the fibre. The Peltier actuators are driven by a DC power amplifier which I built using a power field effect transistor in a push-pull configuration. The temperature of the fibre can be changed by 25°, leading to a maximum optical path length change of 50 mm.

Optical amplifiers The optical amplifiers used for the user and the returned pulse train are manufactured by CzechLight Amplifiers (CLA-Optokon) and they consist of

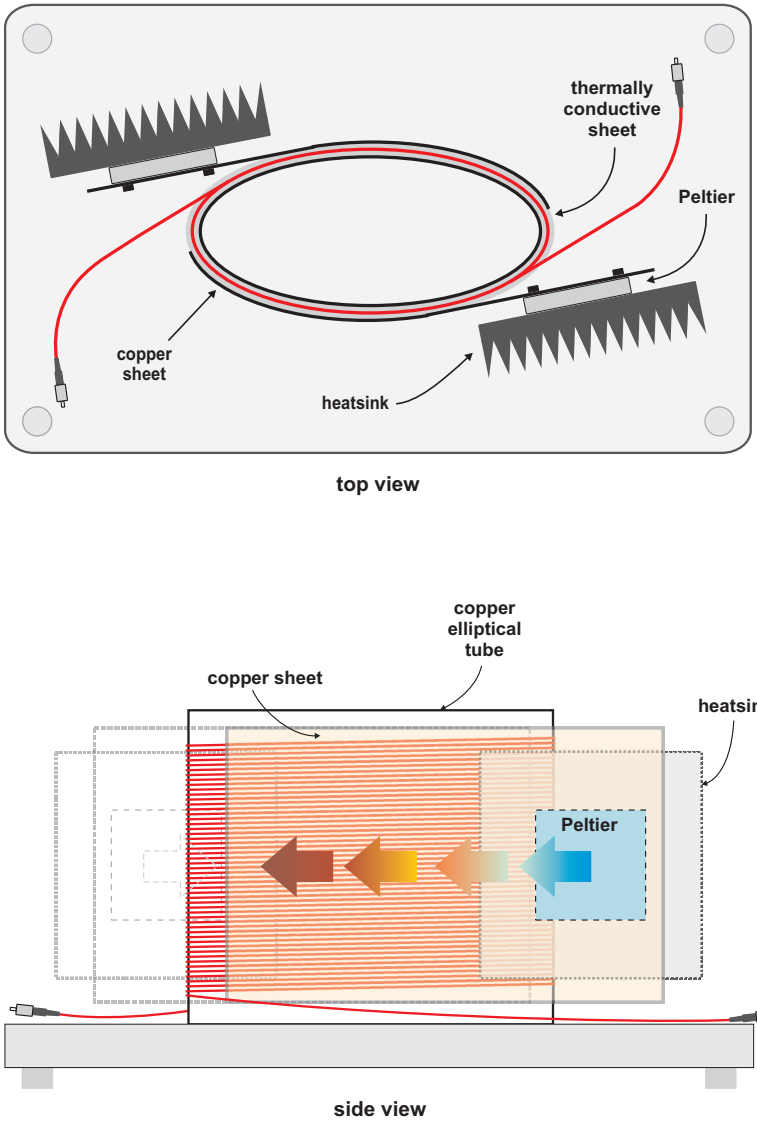

Figure 6.7: Illustration of the design of the thermally controlled fibre spool used to extend to 50 mm the optical path length range of the fibre stretcher. Approximately 300 m of fibre are wound on an elliptical cylinder made from copper sheets. Two Peltier actuators heat/cool the fibre through copper sheets bracing the copper cylinder. Thermally conductive sheets improve the thermal contact between the copper sheets and the fibre. Temperature variations of the order of 25 ° are possible with this arrangement.

two stages, a preamplifier and a booster stage. In this experiment only the preamplifier sections were used. The minimum acceptable input power is −30 dBm and the maximum output power is +15 dBm. The nominal gain is approximately 20 dB and the noise figure is 4.3 dB.

Dispersion compensating modules A selection of dispersion compensating modules were purchased by the ORC to compensate for the dispersion of each section of the

network between the ORC and UCL. A combination of these modules with the addition of a 3.5 km spool of SMF-28 fibre at the remote end allowed the pulses to be recompressed at the remote and returned repetition rate ends to less than 100 ps. The exact length of the fibre is not revealed by the manufacturer but is usually 1/7 of the length of the single mode fibre that the DCF is matched to.

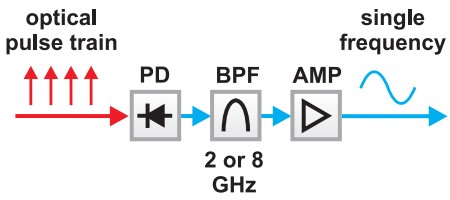

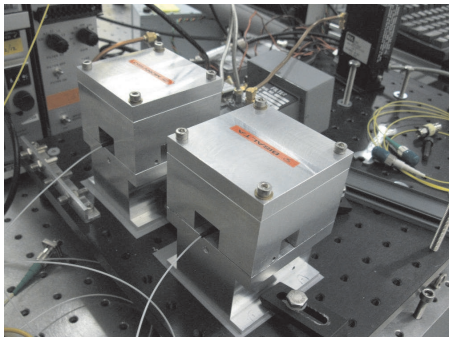

Figure 6.8: Left: block diagram of the repetition rate detection stage, consisting of a photodiode, a narrow band pass filter and an amplification stage. The amplification stage can consist of 3 or 4 cascaded microwave amplifiers depending on the microwave power available at the output of the photodiode. Right: a picture of the fast photodiode enclosure I designed to increase their thermal mass and provide appropriate RF screening.

Detection stages The block diagram of the repetition rate detection scheme used to extract a harmonic of the repetition rate at each end of the fibre is shown in Figure 6.8. Each detection stage consists of a fast photodiode, a narrow band pass microwave filter and a cascade of low noise amplifiers. For the local and remote repetition rate detection, photodiodes with a maximum bandwidth in excess of 12 GHz were used and the 32nd harmonic of the repetition rate (8 GHz) was detected. For the returned repetition rate the detection was performed at 2 GHz as unfortunately a third fast photodetector was not available for this experiment. The fast photodiodes used are from Discovery Semiconductor and I chose them as they are one of the few photodetectors that have been characterized to some extent in terms of AM-to-PM coefficient by the manufacturer and some research groups (Millo et al., 2009; Datta et al., 2009; Zhang et al., 2010, 2011; Taylor et al., 2011). With the aim of substantially increasing their thermal mass for increased temperature stability, I designed an aluminium enclosure for the photodiodes. The enclosures also provide suitable radio frequency screening. Inside the enclosure, two batteries and a current limiter (that I implemented by using simply optoisolator and three resistors) provide a safe bias to the photodiode. The batteries were preferred over an external power supply both for their lower noise and to keep the biasing leads as short as possible to avoid RF pick-up. A picture of the photodiodes in their enclosures is shown in Figure 6.8. If required, two heater resistors can be attached on each side so that the photodiode temperature can either be cycled for testing or stabilized. However, the thermal mass provided by the aluminium enclosures seemed to provide sufficient passive

stability and no more temperature-related characterization appeared to be required for this experiment. The bandpass filters used for the selection of the chosen harmonic allow a suppression of the neighbouring harmonics of more than 30 dB at both 8 GHz and 2 GHz. The microwave amplifiers are from Hittite Inc. Each amplifier is mounted on a large heatsink to avoid temperature-related degradation of the phase noise performance at low offset frequencies from the carrier. Each set of amplifiers was enclosed in a 19 inch aluminium box with SMA feedthroughs. This was a necessary step towards achieving a suitable isolation between the amplifier chains used for the detection of the local and remote repetition rate. As the required amplification exceeded 50 dB, crosstalk between the output of one amplification chain (where the power can be as high as 16 dBm) and the input of the other was particularly easy to observe without taking mitigation steps. It is important for this situation to be avoided as otherwise the power available at the output of the photodetectors can be comparable to the injected power into the amplifiers due to crosstalk.

6.5.2 Reflections

Unlike to the experiment on fibre spools, in this experiment the pulse train propagates from the transmitter to the user and back through the same fibre. Any reflections of the forward pulse train cannot be distinguished from the pulse train returned from the user end at port 3 of CIR1. This determines the minimum optical power that needs to be sent back from CIR2, once the fibre losses are accounted for. A summary of the losses from the output of CIR1 to the input of CIR2 is in Table 6.2

Item	**Attenuation (dB)**
circulator 1, port 1 to 2	0.8
fibre stretchers	0.4
thermally controlled spool	0.5
circulator 2, port 2 to 3	0.8
fibre link	26
interconnections	1
Total	**29.5**

Table 6.2: Optical attenuation budget.

The isolation between port 1 and 3 of circulator CIR1, which is in excess of 55 dB, sets the affordable level of reflections. In order for the returned optical signal to be substantially higher than that reflected by CIR1, say by at least 20 dB, the optical power injected at the user end should be (all powers expressed in dBm)

$$(P_{\mathrm{r}} - 29.5) - (P_{\mathrm{in}} - 55) \geq 20\,\mathrm{dB} \tag{6.5.1}$$

where P_{r} is the power relaunched into the network from the user end and P_{in} is the power injected into the network from the transmitter end. It follows that (all powers expressed in dBm):

$$P_{\mathrm{r}} \geq P_{\mathrm{in}} - 5.5\,\mathrm{dB}. \tag{6.5.2}$$

From the specifications of the optical amplifier used at the remote end, the input power is required to be higher than −30 dBm and for proper low noise operation of the amplifier, it is desirable to have at least −25 dBm input power. This means that the optical power launched into port 1 of CIR1 has to be higher than 5 dBm. It follows that the optical power sent from the user back to the transmitter has to be at least

$$P_{\mathrm{r}} \geq P_{\mathrm{in}} - 5.5\,\mathrm{dB} \approx 0\,\mathrm{dBm} \tag{6.5.3}$$

6.6 Preliminary tests

In preparation for the full frequency transfer experiment I ran a number of preliminary tests, both at NPL and at the ORC. The higher detection frequency required a new SSB modulator to be assembled and characterized as well as the characterization of the performance of the new fast photodiodes. These tests were performed at NPL. At the ORC I first measured the phase noise of the fibre link which provided me with useful information about the required optical path compensation range whilst planning the experiment set-up. I also compared the measured phase noise to that measured on other networks in other countries and to that measured on fibre spools. These tests are described in the following paragraphs.

6.6.1 Single sideband modulator at 8 GHz

The SSB modulator is used to shift the frequency in one of the two measurement arms (local or remote repetition rate) allowing a direct frequency stability measurement to be taken with a frequency counter. In contrast to the previous experiment where a 1.5 GHz SSB modulator constructed from discrete components, this time I used an commercial SSB modulator. The smaller package provides improved temperature balance between the two arms of the modulator, resulting in a better performance at long time scales. As for the previous SSB modulator, a carrier suppression arm was added in order to reject most of the unmodulated signal. The whole modulator was mounted on an aluminium

breadboard and enclosed in a 19 inch aluminium box. The performance of the new 8 GHz SSB modulator, measured in the same way as for the 1.5 GHz modulator, is shown in Figure 6.9. For comparison the 1.5 GHz SSB performance is also shown. The substantial improvement is due to the increased frequency and, at long term, to the improved thermal balance between the two arms of the modulator, as discussed above.

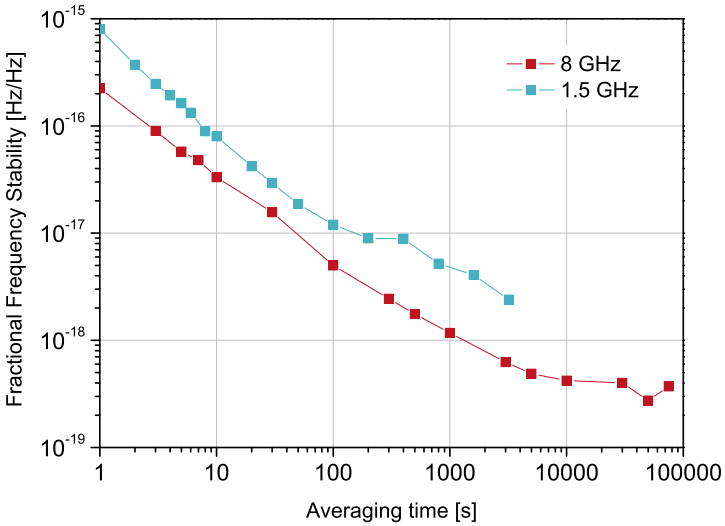

Figure 6.9: Improvement in the stability of the SSB modulator when the frequency is increased from 1.5 GHz to 8 GHz.

6.6.2 Phase noise of the 86 km-long installed fibre link without noise cancellation

The phase noise induced on the repetition rate by the acoustic noise and thermal effects in the fibre was measured using the set-up shown in Figure 6.10. The 8^{th} harmonic of the repetition rate is detected before and after the 86 km link. The extracted repetition rate is then phase compared on a microwave mixer with the inputs set in quadrature via a manual phase shifter.

The measured phase noise is shown in Figure 6.11. For comparison, the phase noise measured on fibre spools (100 km) at NPL is also shown. A higher level of low frequency noise is found for the installed fibre case, despite the fibre being 14 km shorter. The higher noise could be attributed to thermal effects resulting from the installed fibre being exposed to temperature changes larger than those typical of a laboratory environment. Also the noise above 1 kHz is higher. Interestingly, the phase noise between 50 Hz and 500 Hz is lower for the installed fibre than the spooled fibre. This might be a confirmation of the phase "kick" build-up effect described at the beginning of this chapter resulting from the geometry of the spooled fibre. In the inset plot, the bumps due to self-heterodyning of the laser noise are shown. The dips are at multiples of approximately 2

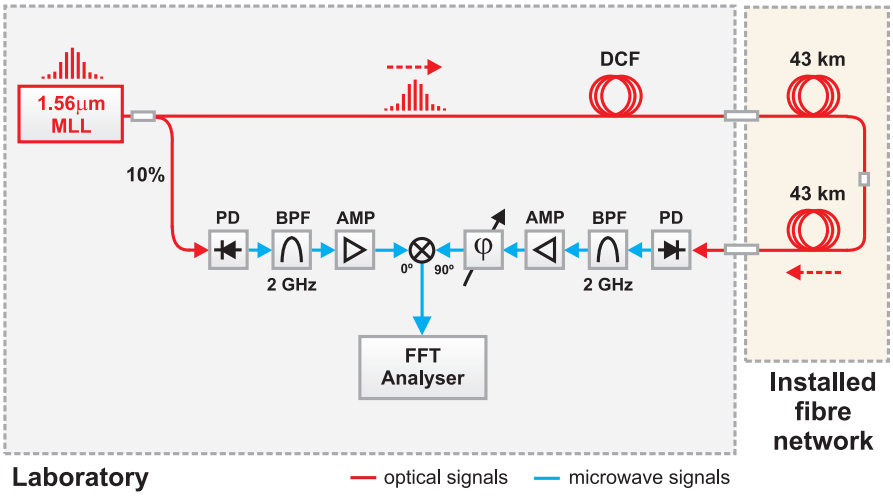

Figure 6.10: Set-up for the measurement of the phase noise of the 86 km-long JANET-Arora link.

kHz. This frequency agrees very well with that calculated from the propagation time of the pulse train over the full length of fibre. The length of the DCF module is estimated to be approximately 12.3 km, from which follows

$$f_{\mathrm{dip}} = \frac{v_{\mathrm{fibre}}}{L_{\mathrm{SMF}} + L_{\mathrm{DCF}}} = \frac{2 \times 10^8\,\mathrm{m}}{(86 + 12.3) \times 10^3\,\mathrm{m \cdot s}} = 2034\,\mathrm{Hz} \tag{6.6.1}$$

where v_{fibre} is the speed of light in an optical fibre (approximated to 2×10^8 m s^{-1}), L_{SMF} and L_{DCF} are the length of the SMF and DCF fibres respectively.

The noise at offset frequencies lower than 10 Hz appears to follow a $1/f^2$ trend, although a better estimation of the correct slope could be derived only from the Fourier analysis of longer data sets. Rather than setting the FFT analyser to lower offset frequency, I preferred to acquire the phase changes with a digital voltmeter and then perform the FFT on the time domain data and two samples of this data are shown in Figure 6.12 over time scales of up to 2000 seconds. The reason for preferring this approach is that if phase glitches occur during the computation of the spectrum by the FFT analyser it takes a long time for the perturbation to average out and for the measurement to be trustable again. If the phase noise information is instead acquired in the time domain, any glitches can be easily identified and the appropriate glitch-free part of the data can be selected for analysis. From Figure 6.12 it can be seen that phase excursions on the installed fibre are not only larger, but also much more irregular than the almost linear drift experienced on fibre spools at NPL. This can be explained by observing that the temperature evolution of a fibre spool is likely to be much more homogeneous than that of an installed fibre. Phase changes for the installed fibre were measured to be

as large as 1.1 rad in less than half an hour. The spectrum calculated from the time domain phase changes is also shown in Figure 6.12. Here, the noise slope can be more reliably approximated and is confirmed to be proportional to $1/f^2$. The noise process corresponding to this slope is a random walk of phase. This observed noise process can be explained by noting that temperature effects usually do have a random walk-like type of behaviour. The phase is in fact linearly proportional to the optical path length, which in turn is linearly proportional to the temperature, leading to the phase changes mimicking the random walk behaviour of the temperature changes.

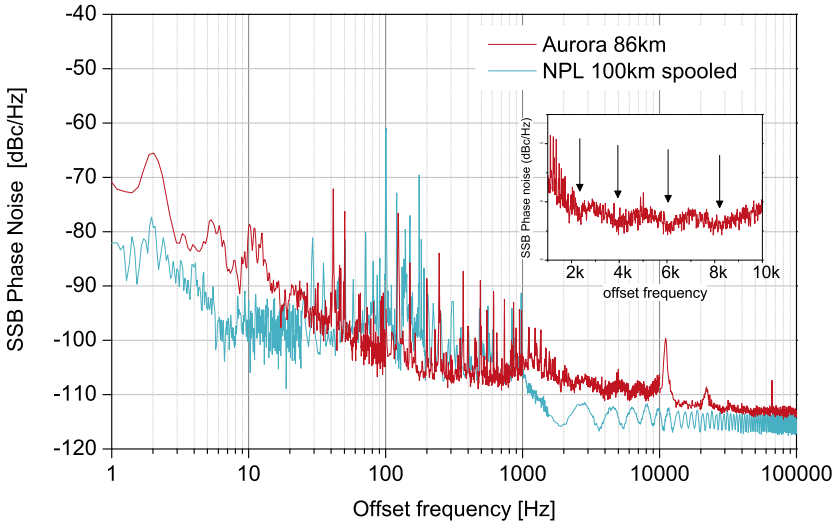

Figure 6.11: Phase noise of the 86 km-long JANET-Aurora fibre link between the ORC at the University of Southampton and the hub at Crawley Court. The noise is higher for low offset frequencies, compared to that of a fibre spool in a temperature controlled laboratory environment, because of the the higher temperature variations of an installed fibre. The acoustic noise between 50 and 500 Hz is, in contrast,lower. Inset, the perioc structure of the phase noise arising fron self-heterodyne detection of the laser noise for offset frequencies higher than 1 kHz is visible.

6.6.3 Comparing the Aurora network to other networks in other countries

To my knowledge, no other phase noise data measured by propagation of an optical comb over an optical fibre link of tens of km is available at the time of writing. The only measurements available are those performed on the 3.45 km-long fibre link between JILA and NIST. However, it is possible to convert the phase noise measured by propagating a frequency comb to that which would be measured when a single optical carrier is propagated (see chapter 3). In Figure 6.13, the phase noise measured on the 8th harmonic of the repetition rate (2 GHz) has been upconverted to optical frequencies.

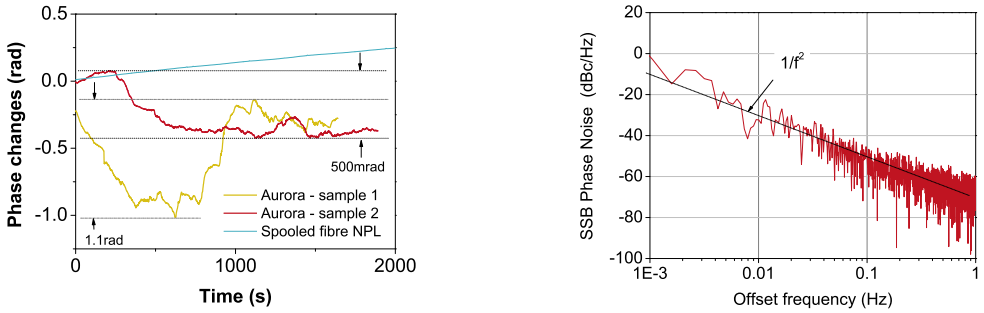

Figure 6.12: Time and frequency domain representation of the measured phase noise on the JANET-Aurora link. The slope of the phase noise spectrum between 1 mHz and 1 Hz indicates that the noise corresponds to flicker of frequency (random walk of phase).

For comparison, the noise measured by other groups using a single optical carrier is shown. Up to 100 Hz offset the resulting phase noise is comparable to those measured on links of similar length at PTB and LPL-SYRTE. The noise measured at NMIJ is substantially higher. This should be expected to some extent as, unlike the PTB and LPL-SYRTE cases, the fibre link runs across metropolitan areas for most of its length. In the measurement with the comb, for offset frequencies higher than 100 Hz, the self hetherodyne noise of the laser and detection noise dominates that of the fibre and no comparison can be made with the phase noise measured on the optical links in other countries.

6.7 Experimental results

In this section the experimental results for the transfer of a microwave reference over the Aurora network using an optical frequency comb are described. The indicators of the performance of the transfer technique are the frequency stability and the phase noise spectrum, which will be found at the end of this chapter. These quantities were measured also for the previous experiment on fibre spools and therefore a comparison can be made.

6.7.1 Quality of the detected repetition rate

The first step towards the cancellation of the phase noise introduced by the fibre is the optimization of the detected harmonics of the repetition rate at the local and remote end. The optimization involves adjustments of the optical powers at different points of the experimental scheme such as the power launched into the network and the output power of the optical amplifiers. The main concern is the power ratio between the powers of the returned pulse train and the reflections at port 3 of CIR1 (6.6). The optical

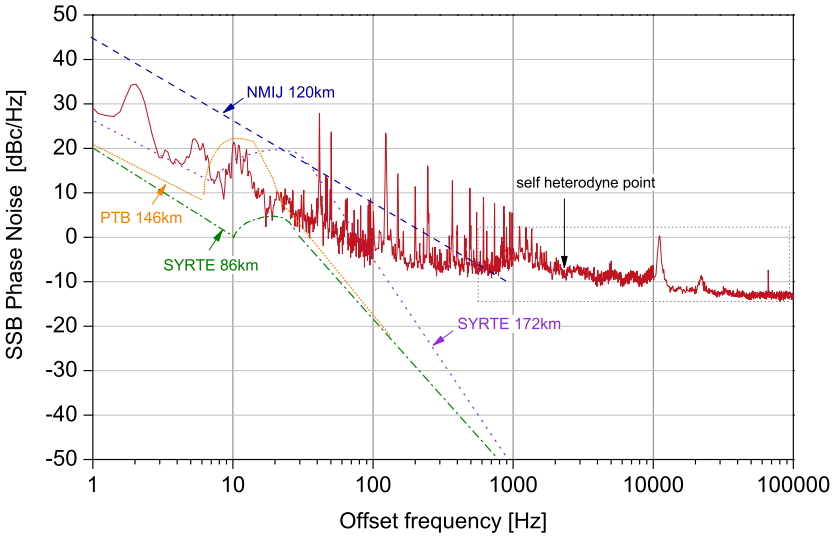

Figure 6.13: Comparison between the phase noise measured on a number of fibre links in different countries to that measured over the JANET-Aurora link. The phase noise was measured directly at optical frequencies at PTB, LPL-SYRTE and NMIJ, whilst it has been scaled up to optical frequencies in the case of the NPL-ORC comb transfer experiment as it is measured at a microwave frequency (2 GHz).

spectrum of a typical ratio between these two powers, after optimization, is shown in Figure 6.14 and is observed to be in the region of 15 dB. The returned optical pulse train exhibits a narrower bandwidth as it is limited by the optical amplifiers. The part of the spectrum of the reflected pulse train with wavelengths higher than 1570 nm does contribute to the unwanted detection of the repetition rate at port 3 of CIR1. A reduction of the level of these reflections could be achieved by inserting a 30 nm filter at the input of the experimental setup as it would reject those optical frequencies that are outside the bandwidth of the optical amplifier. However, the reflection levels appear sufficiently low not to be the limiting factor to the transfer stability performance. The dispersion compensating modules, with the fine tuning provided by the additional 3.5 km-long fibre spool, allowed for the 32nd (8 GHz) and 8th (2 GHz) harmonic of the remote and returned repetition rate respectively to be extracted with a signal-to-noise ratio in excess of 70 dB on a 1 kHz resolution bandwidth.

6.7.2 Amplitude noise

Amplitude noise was found on the remote and returned repetition rate and is shown as the grey line in Figure 6.15. The amplitude noise was measured using a microwave power detector. The amplitude noise is a particular problem for the phase comparison between the local and returned repetition rates at the transmitter end. The mixer used

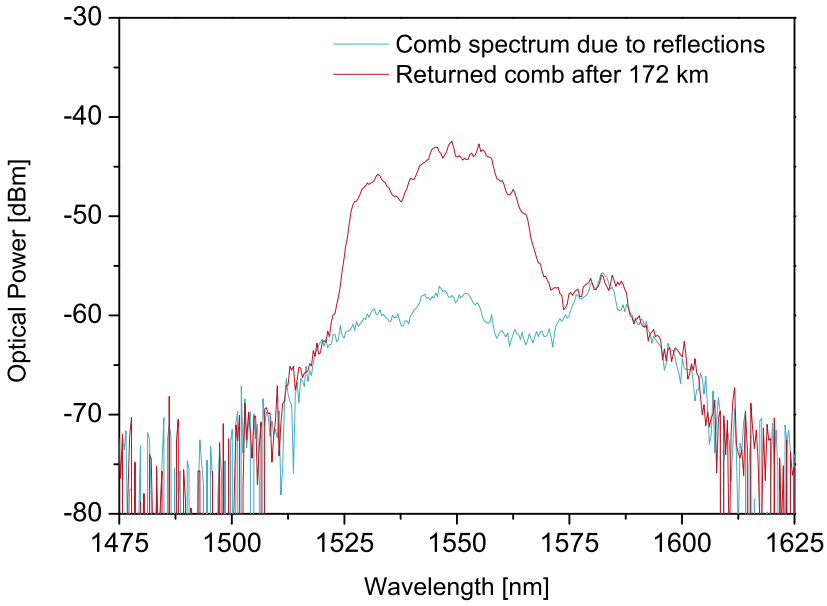

Figure 6.14: Optical spectrum of the returned and reflected pulse train as detected at port 3 of the circulator CIR1.

for the phase comparison has a non-zero sensitivity to changes in the power applied to its input ports. Changes of the output voltage of the mixer due to amplitude changes of the input signals cannot be distinguished from those due to phase changes. The output voltage will be integrated and a correction will be applied to the optical path length of the fibre link via the fibre stretcher, effectively converting amplitude noise into phase noise. The sensitivity of the mixer to amplitude variation depends on the phase relationship between the input signals. Most mixers exhibit a minimum of the AM sensitivity reasonably close to, but not exactly, at the quadrature condition (90 degree phase shift between the inputs). The output voltage of the mixer for various quasi-quadrature points is shown in Figure 6.15 when AM noise is present on one of the two inputs. The red trace shows the output voltage when the mixer is operated at the "'sweet spot". Unfortunately, in the comb transfer experiment it is not possible to preserve this condition for long times as the phase variations of the input signals are quite large, so tuning the mixer onto the "sweet spot" is not a viable solution. In order to avoid AM-to-PM conversion, the bandwidth of the fibre phase noise cancellation loop was limited to approximately 10 Hz. Excess amplitude noise was also observed after the fibre when the input spectrum presented narrow features ("spikes"'). I attributed the excess noise to non linear effects taking place in the fibre. By adjusting the amplifier squeezers (see chap. 4) it was always possible to smooth out the spectrum and remove the narrow features and, consequently, also the excess amplitude noise.

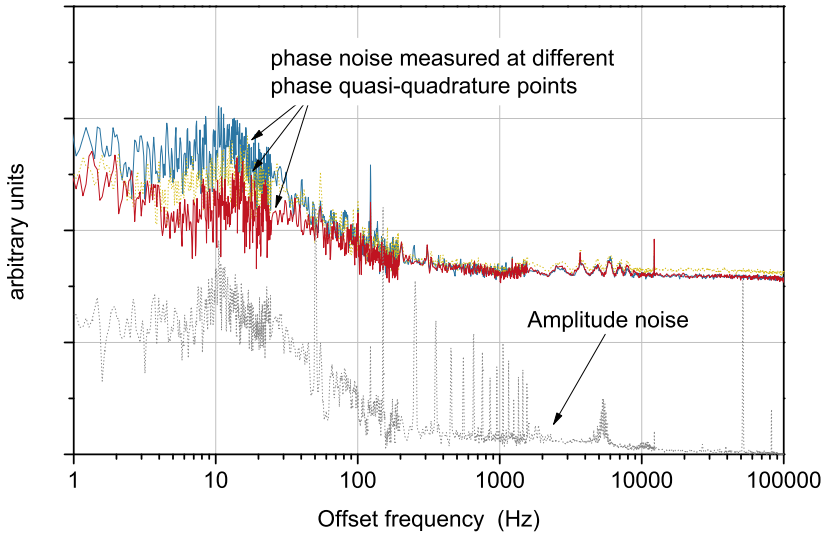

Figure 6.15: Phase noise measured with different quasi-quadrature conditions. The grey trace is the measured AM noise on the repetion rate with an RF detector.

6.7.3 Measurement of the passive transfer stability

The transfer stability achievable without any suppression of the noise introduced by the fibre is shown in in Figure 6.16. I have shown the free running transfer stability for the experiment on fibre spools at NPL as a comparison to provide a further visual explanation of the differences between the two cases. The more uniform temperature distribution along the fibre in the spooled fibre case clearly shows up in the stability plot when compared to the installed fibre case. The short term stability, for time scales of a few seconds, does not differ substantially between the two cases. On these time scales, thermal effects play a smaller role and the substantially different environments in which an installed fibre and fibre spools in a quiet laboratory are located, show themselves to a lesser degree. It is interesting to note that even in the installed fibre case, the measured transfer stability is still better than 10^{-14} for time scales greater than 5 s.

6.7.4 First transfer stability measurement with the fibre phase noise cancelled

The result of the first measurement of the transfer stability over 86 km of installed fibre is shown in Figure 6.17. A clear improvement of the stability delivered at the user end for time scales greater than 10 seconds is visible. However, the short term stability (<10 s) is only marginally better. By repeating the frequency stability measurement under different conditions, such as different driving current to the optical amplifiers (from

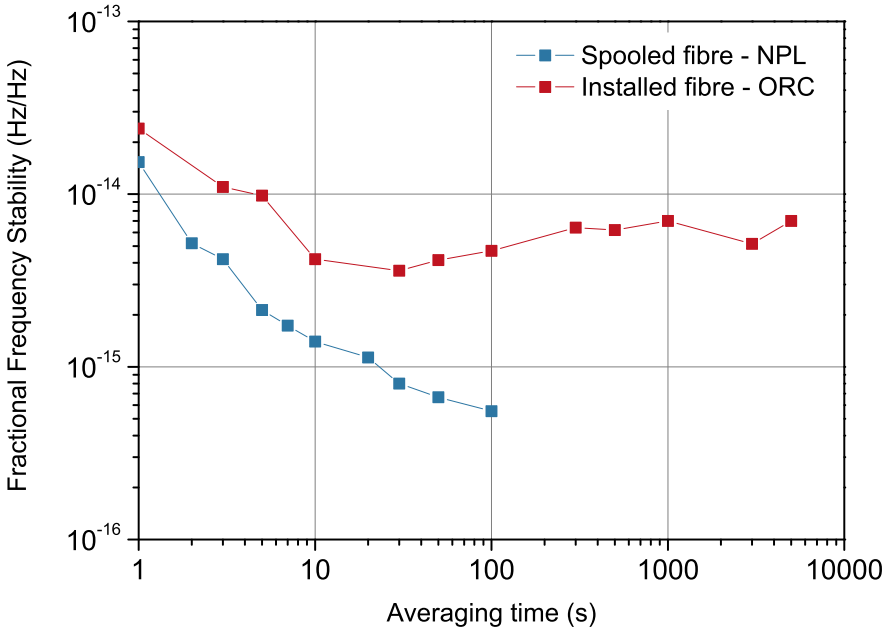

Figure 6.16: Fractional frequency stability of the repetition rate detected at the user end when no fibre noise cancellation is performed. As expected, the stability is worse for the installed fibre case, both because of its increased length and because of the higher temperature fluctuations experienced by the fibre. Interestingly, the stability is still better than 10^{-14} for most time scales.

Optokon-CzechLight), it became evident that the amplifiers played some role in limiting the level of short term stability achievable. I then tested the amplifiers separately by comparing the repetition rate detected at the output of the optical amplifier to that detected at the input. Any noise produced by the amplifier will be detected in this way, independent of the noise of the laser source. The optical amplifiers showed a peculiar behaviour. The measured stability appeared to be strongly dependent on the driving current and even a small change in this current could cause a large difference in stability. A typical situation is shown in Figure 6.18 where a change in the current as small as 5 mA (corresponding to a change in the output power of a fraction of a dB) was sufficient to change the measured stability by up to a factor of four at 1 s. No appreciable changes in the output optical spectrum were observed when the current was changed by these amounts, excluding to first order the possibility of nonlinear effects taking place. This peculiar behaviour was observed in other amplifiers of same model. The test was then repeated using an optical amplifier manufactured by Keopsys. In this case the stability was measured to be consistent over a large range of driving currents and also a better level of stability was achieved. The achieved stability degraded only for very low input power levels as one might expect. The results obtained with this latter amplifier is shown on the right of Figure 6.18. Unfortunately, only one Keopsys amplifier was available for this experiment, whilst two are required for the user and returned repetition rate detection. Many months after the completion of the fibre transfer experiment on

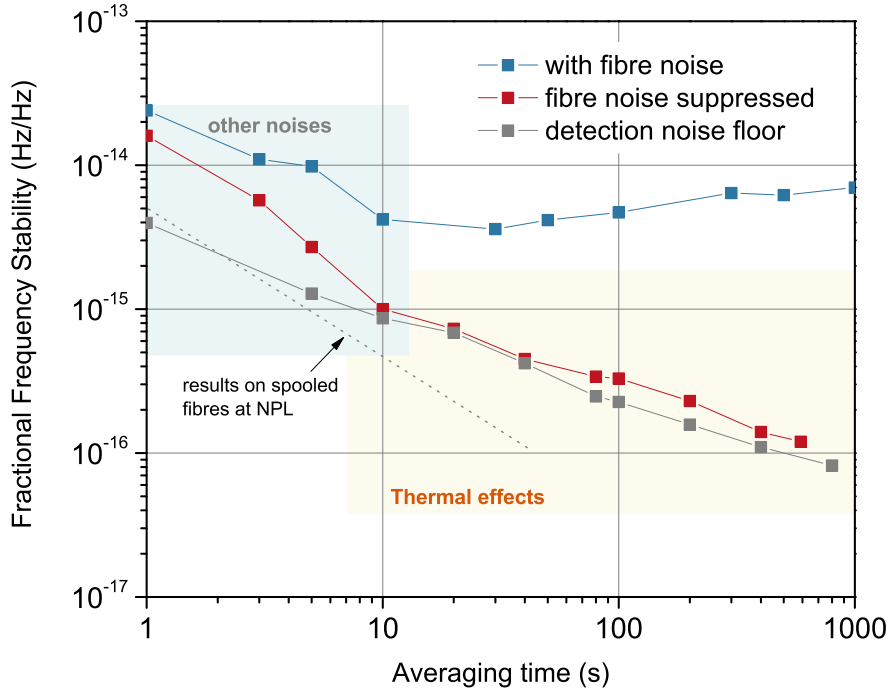

Figure 6.17: Fractional frequency stability of the 32nd harmonic of the repetition rate detected at the user end when the fibre noise is cancelled. The short term stability is only marginally better than when the noise is not cancelled and does not match the measurement noise floor, indicating that there might be a source of excess flicker phase noise to be identified.

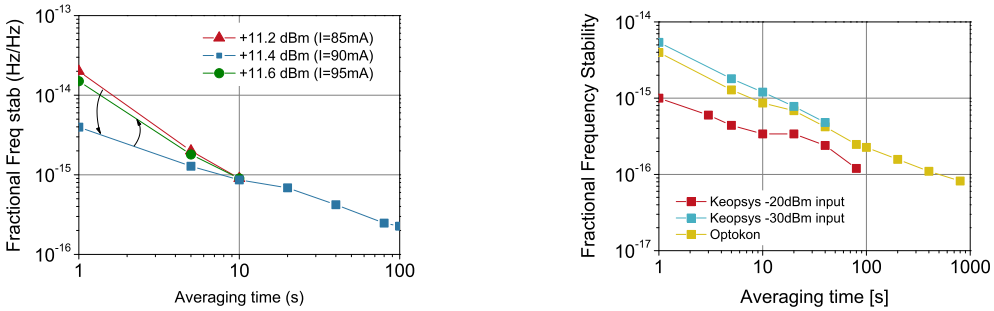

Figure 6.18: Left: residual fractional frequency stability of the Optokon-CzechLight amplifier under different driving current conditions. Even for current changes as small as 5 mA their performance can be substantially different. Right: the better stability achieved by the Keopsys amplifier.

the Aurora network, the manufacturer has identified and resolved a problem with the amplifier's control electronics which was responsible for the excess noise.

6.7.5 Final results

Here I report the final microwave transfer results achieved when the CLA amplifiers are operated with a suitable drive current.

Transfer frequency stability The transfer frequency stability is shown in Figure 6.19 and is 5×10^{-15} at 1 s, reaching 4×10^{-17} at 1600 s. The measured transfer stability over the Aurora network matches that measured when a short piece of fibre is used instead. For long time scales, this suggests that possible detrimental effects, such a polarization mode dispersion (PMD), do not yet play a relevant role at the measured level of stability and over the ime scales tested so far. PMD was found to be a limiting factor in microwave frequency transfer experiments performed by LPL-SYRTE by intensity modulation of an optical carrier at the 10^{-16} level, particularly for time scales of the order of 1000 s (Lopez et al., 2008, 2010). The PMD is a possible source of instability because causes the propagation delay to be different for the forward and backward direction, breaking the symmetry required by the phase noise cancellation principle. However, in the LPL-SYRTE experiments the PMD problem was enhanced by the fact that two lasers at different wavelengths and with different polarization states were used for the forward and backward direction of propagation. In the comb experiment the same optical signal is returned to the transmitter. It is possible however that PMD effects could be observed for much longer time scales or for much longer links. The short term stability also matches that measured at NPL over 50 km of spooled fibre, shown by the green dashed line. The free running stability is also shown for comparison, showing that an improvement of more than two orders of magnitude is achieved for time scales longer than 1000 s. At the time of writing of this thesis the demonstrated stability is the highest ever reported for a microwave frequency transfer based on the propagation of an optical frequency comb and using phase detection techniques in the microwave domain.

Phase noise The phase noise plot is shown in Figure 6.20 and is observed to be suppressed by up to 15 dB for offset frequencies below 10 Hz when the feedback loop is activated. At 1 Hz offset from the carrier the single sideband phase noise is measured to be -84 dBc/Hz. The dips corresponding to multiples of the inverse of the transit time through the fibre (SMF+DCF) can be observed.

Jitter It is interesting to convert both the measured frequency instability and phase noise into timing jitter. One of the possible applications of this technique is the distribution of ultra precise timing signals, for example in particle accelerators or arrays of radiotelescopes, and the average jitter of the optical and microwave pulses is the quantity of interest for these communities. From the frequency stability value for each averaging time it is possible to calculate the corresponding timing jitter over that time scale as

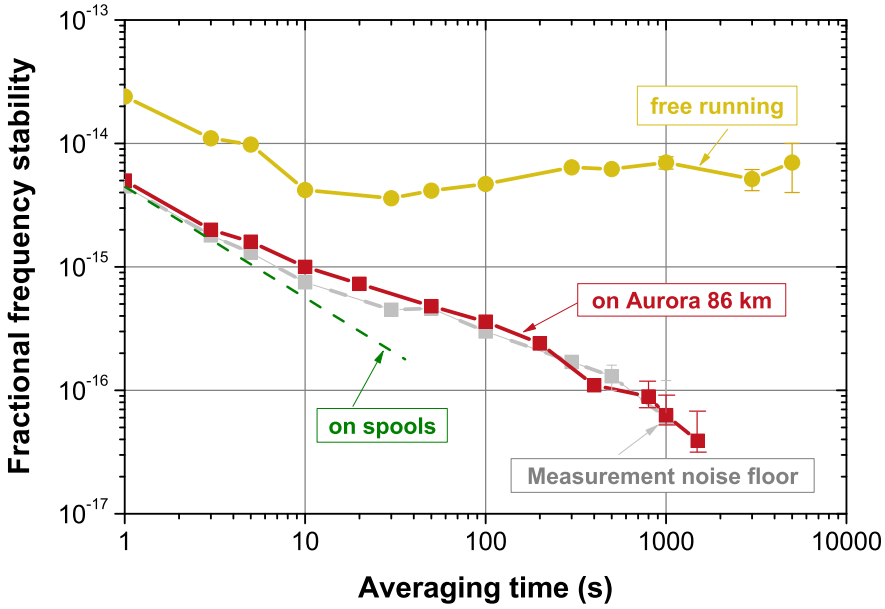

Figure 6.19: Final fractional frequency stability results. The 32nd harmonic of the repetition rate is detected at the user end when the fibre noise is cancelled. The short term stability now coincides with that measured when the fibre link is replaced by a short fibre. For comparison, the previous results achieved on 50 km of spooled fibre in the NPL laboratory are shown. At short term the results on installed fibre overlap with those achieved on spools. For longer time scales the stability for the installed fibre case is worse because of the excess noise introduced by the optical amplifiers.

$$T_{\mathrm{rms}} = \frac{\delta f}{f} \cdot \tau \tag{6.7.1}$$

where T_{rms} is the timing jitter (rms value), $\frac{\delta f}{f}$ is the fractional frequency stability and τ is the averaging time. At 1600 s the measured stability is 4×10^{-17} and so

$$T_{\mathrm{rms}} = 4 \times 10^{-17} \cdot 1600\,\mathrm{s} = 64\,\mathrm{fs}. \tag{6.7.2}$$

From the phase noise measurement it is possible to calculate the timing jitter over much shorter time scales. Integrating the phase noise from 0.1 Hz to 100 kHz (10 μs to 10 s) the timing jitter is 82 fs. However, integrating from 1 kHz to 100 kHz gives a timing jitter of 80 fs, showing that most of the jitter is due to the bumps in the phase noise beyond a few hundred Hz. These bumps arise from the self-heterodyne effect and their magnitude depend on the gain setting of the repetition rate locking, which could be set

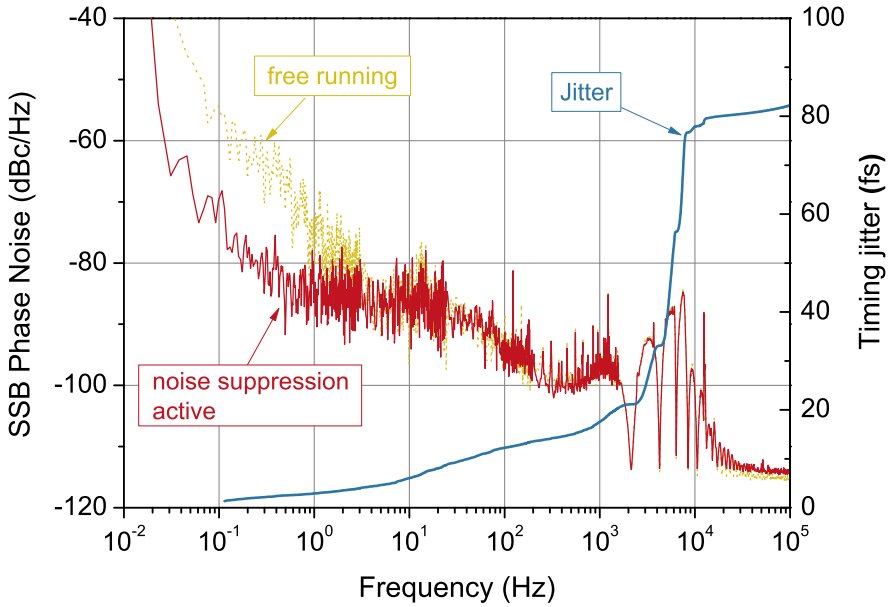

Figure 6.20: Phase noise of the 32nd harmonic of the repetition rate (8 GHz) when the phase noise cancellation loop is switched on (red trace) and off (gold trace). Up to 15 dB suppression is achieved for offset frequencies lower than 10 Hz. The blue trace is the total jitter and is measured to be 82 fs when the phase noise is integrated from 1 Hz to 100 kHz.

to lower values than in this particular measurement[3]. In addition, one should notice that the noise actually present on the repetition rate at the user end will be 6 dB lower (and no dips will be present) as the formula of the self-heterodyne effect transfer function contains a factor of 4 (see Equation 4.2.4).

Accuracy A crucial measurement of the performance of the transferred frequency is the accuracy test. It is possible that the frequency delivered at the user end of the fibre is very stable but that it is affected by a systematic frequency offset. It is therefore necessary to compare the beat frequency produced by mixing the repetition rate delivered at the user end with that frequency shifted by the SSB modulator at the local end with the frequency produced by the DDS that drives the SSB modulator. The difference between the mean value of these frequencies measured over a certain time scale should be within the error bars of the stability measured over the same time scale. This condition is verified as the mean fractional frequency difference was measured to be 1.4×10^{-17} over 1600 s.

[3]The increase in the repetition rate noise for offset frequencies higher than 1 kHz is due to the piezo controlling the cavity length in the laser approaching resonance. The phase noise overshoot at high offset frequencies shown in Figure 6.20 was observed only after taking the frequency stability measurement of Figure 6.19 therefore, although smaller noise bumps could be achieved at lower gain settings, the gain was left the same to maintain consistency of experimental conditions.

6.8 Comparison of the results with other groups

In Figure 6.21 the comb transfer stability results achieved over 86 km of the JANET-Aurora network are compared to those achieved by Holman (Holman et al., 2005) on the 3.5 km-long installed fibre link that joins JILA to NIST. The transfer stability achieved over the Aurora fibre link is better by more than a factor of two for most time scales and it is the highest ever reported for a microwave transfer technique employing optical combs propagated over installed fibre, at the time of writing this thesis.

Three novel aspect characterize the results:

1. measurement of the true transfer stability delivered at the user end;
2. different phase noise cancellation scheme;
3. transfer stability measured for time scales greater than 100 s for the first time.

Demonstration of the true transfer stability delivered at the user end In the JILA experiment, no test of the frequency stability delivered at the user end (at the midpoint located at 3.45 km away in the NIST laboratory) was performed. A separate set of photodetectors and amplifiers were used for the measurement of the frequency stability of the repetition rate delivered after 6.9 km. However, this arrangement would detect only the possible extra noise due to the detection stages, not instabilities of the repetition rate that might be present at the user end, resulting effectively as an in-loop measurement. The stability at the user end can be trusted to be the same as that measured after a round trip only under the assumption that the noise is sufficiently correlated between the two fibres. In the experiment over the Aurora network the stability effectively delivered to the user end is directly measured and no such assumption is necessary.

Different phase noise cancellation scheme In the JILA experiment, the pulse train was propagated in a unidirectional fashion through the fibre. From the mid-point at NIST, the signal was returned via a separate fibre. In the experiment over the Aurora network, the signal was returned through the same fibre for phase noise cancellation.

Transfer stability measured for τ >100 s In the previous experiments, both that performed at JILA and that performed on spooled fibre at NPL, the transfer stability was measured only to a maximum averaging time of 100 s. In the case of the spooled fibre experiment at NPL the measurement of the stability at time scales longer than approximately 60 s was made impossible by the limited range of the fibre stretcher, which was unable to compensate for the long term optical path length variations of 100

km of fibre. The measurement of the transfer stability beyond 100 s is important as it can demonstrate whether second-order effects, such as slow changes in the polarization of the optical signal, pose a limitation to the transfer technique. The comb transfer results on the Aurora network confirm that these second-order effects are not the limiting factor up to the stability levels demonstrated in the experiment. The stability measured over the 86 km-long link when its noise is cancelled is in fact the same as that measured when only a few meters of fibre are used.

In Figure 6.21, the transfer stability measured by LPL-SYRTE over an 86 km-long installed network using an amplitude modulation technique on a single optical carrier is shown for comparison. The microwave transfer stability is up to a factor of 5 to 10 better than that measured when an optical frequency comb is transmitted. However, the comb-based technique delivers to the user not only a comb of ultra-stable microwave frequencies, but also a comb of optical frequencies (15,000 in the experiment described here) where the mode spacing is preserved to a level up to a few parts in 10^{17}. By contrast, the optical carrier in the AM carrier technique is exclusively a (single frequency) microwave transfer technique and the optical carrier delivered at the user end cannot be used as a stable reference as it is only its modulation that is stabilized. The optical carrier could be potentially orders of magnitude more unstable than the microwave frequency that is being transmitted.

6.9 Conclusions

I have transferred an optical frequency comb over 86 km of the JANET-Aurora network and demonstrated that the repetition rate, or equivalently the spacing between the optical frequencies, can be preserved to extremely high levels of stability and accuracy. These results can be interpreted both in the frequency and time domain. The interpretation in the frequency domain gives an indication of whether this technique is suitable for frequency metrology applications. That in the time domain is helpful to asses the technique for precise timing synchronization applications, such as those used in particle accelerators and arrays of radiotelescopes.

Frequency domain The experiment demonstrates that, using an optical frequency comb, it is possible to transfer a microwave frequency with an accuracy and stability of a few parts in 10^{17} over many tens of kilometers of installed optical network. This accuracy and stability level is better than most state-of-the-art microwave and optical frequency references.

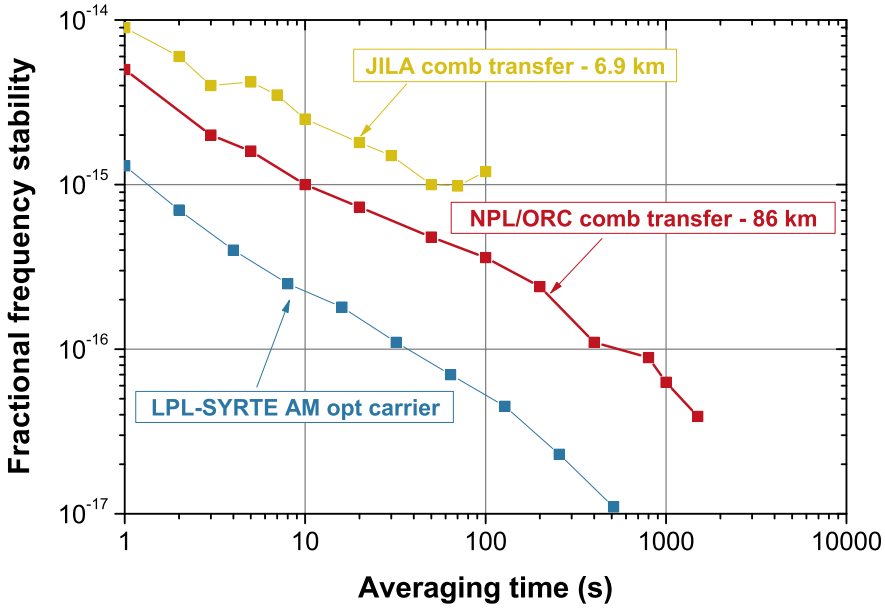

Figure 6.21: Final results showing the fractional frequency stability of the repetition rate detected at the user end of the JANET-Aurora fibre link. The stability reaches 4×10^{-15} at 1600 s. The short term stability has improved from 1.5×10^{-14} to 4.6×10^{-15}. Results obtained with the comb transfer technique by JILA are shown for comparison as well as those achieved using a modulated optical carrier (AM-carrier) technique at LPL-SYRTE. In terms of the stability achieved, the AM-carrier technique is still better for the transfer of a single microwave carrier (marginally when a 1 GHz carrier is transferred (Lopez, 2008), but more substantially when a 9 GHz is used (Lopez, 2010). However, the comb transfer technique can transfer tens of thousands of equally spaced optical modes and hundreds of equally spaced microwave frequencies. Results comparable to those achieved at LPL-SYRTE could be possible if lower noise optical amplifiers were employed.

Time domain The long term timing jitter of the pulse train delivered at the user end is as small as 64 fs over 1600 s, whilst the fast jitter is 82 fs over a 100 kHz bandwidth. These timing jitter levels are compatible with applications where very demanding synchronization is required between two remote locations, such as particle accelerators.

Chapter 7

Optical phase detection and test of the full comb structure

In the previous two chapters, I have described two experiments on the transfer of a microwave frequency over fibre by propagation of an optical frequency comb. The achieved results demonstrated that the technique is capable of transfer stability and accuracy comparable with most of the best frequency references available at the time of writing. The phase detection required for the measurement and the cancellation of the noise introduced by the fibre link was performed in the microwave domain, after optical-to-electrical conversion of the optical pulse train. In this chapter I will discuss how the detection sensitivity can be increased by a factor up of 10^6 when the fibre noise detection is performed in the optical domain rather than in the microwave domain. More importantly, I will present results of the measurement of the stability and accuracy of the individual optical modes transferred to the user end rather than just their spacing. To my knowledge, this is the first time such measurements have been made for a noise cancelled optical link.

The results of this experiment have resulted in the paper "Dissemination of an optical frequency comb over fibre with 3×10^{-18} fractional accuracy" published in Optics Express in 2012 , Marra et al. (2012).

7.1 Preserving the full structure of a transmitted optical frequency comb

A frequency comb is described by two parameters, the tooth frequency spacing (mode spacing) and the carrier envelope offset frequency. The mode frequency spacing can be detected by optical-to-electrical conversion with a photodiode and the experiment described in the previous chapters showed that it can be preserved to a very high accuracy

after transfer over many tens of kilometer of fibre. However, the whole comb could still "wander" in frequency whilst preserving its tooth spacing. This situation is depicted in figure 7.1. If the chromatic dispersion changes over time, the relationship between the envelope of the pulse and the phase of the optical carrier must change, leading to the whole frequency comb shifting in frequency without changing the frequency spacing between the optical modes. In order to assess to what level the entire structure of the comb is preserved when transmitted over fibre, the stability and accuracy of the individual optical modes delivered at the user end must be measured in addition to the mode spacing.

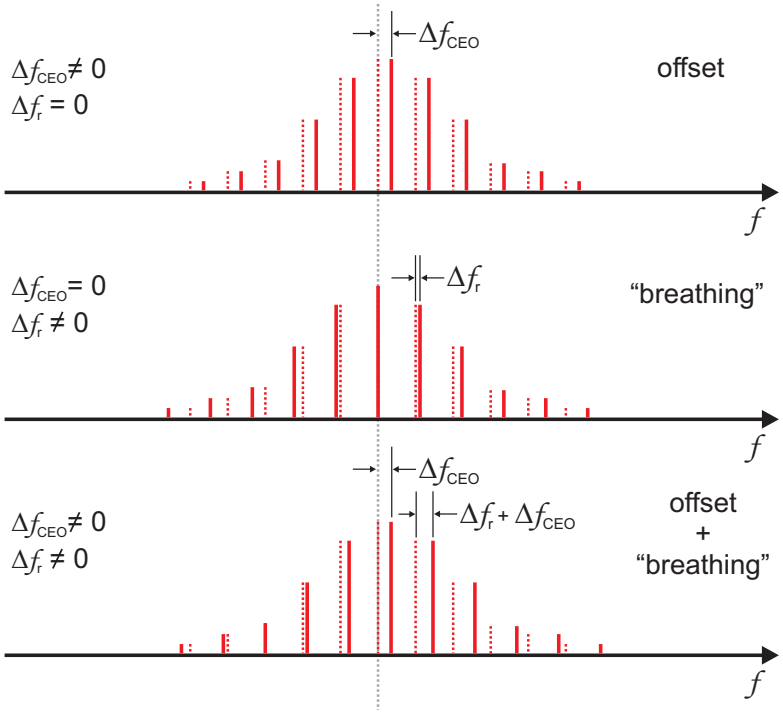

Figure 7.1: The frequency of the optical modes in an optical frequency comb depends on two parameters: the carrier envelope offset frequency f_{ceo} and the repetition rate f_r. When the repetition rate changes, the optical modes whose frequency is higher than the mode at the center of the comb (for an ideal comb) will increase their frequency, whilst those with frequency lower than the center mode will decrease it. This effect is often called "breathing". In a frequency transfer experiment the dispersion of the fibre could change over time and the carrier envelope offset frequency will therefore change.

7.2 Principle of the optical phase detection

The measurement of the optical path length changes in the fibre by detection of the Doppler-induced changes of the repetition rate is limited by the noise introduced during the optical-to-microwave conversion with a photodiode and the microwave devices used to condition the converted signal. Much higher sensitivity to the optical path

length changes can be achieved if phase changes due to the fibre noise are measured in the optical domain. The technique described below allows exactly that, increasing the phase sensitivity by a factor of up to 10^5. The idea behind the technique came to me after asking myself whether it would have been possible to transpose the interferometric techniques used in the single carrier transfer technique onto thousands of optical comb modes. I later discovered that the technique is similar to one of the two techniques used previously by Chen et al. (2006) for using combs to transfer of a microwave signal over 60 m of fibre. In that experiment however, no investigation of the stability of the optical comb modes was undertaken. Also, the phase noise of the repetition rate was effectively an in-loop measurement after a round trip through the fibre rather than the phase noise delivered at the "user" end.

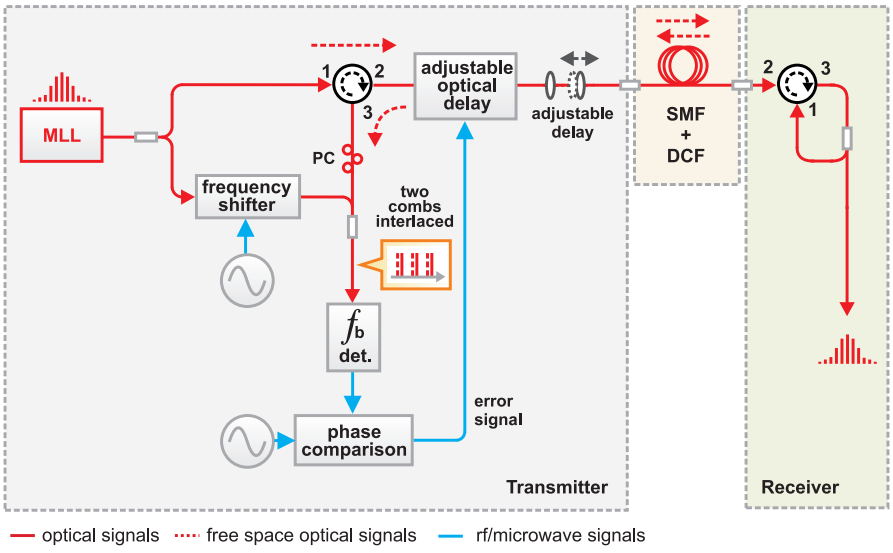

Figure 7.2: Block diagram illustrating the optical phase detection principle. The optical frequency comb is sent through the fibre and back and then it is recombined with the optical signal available at the output of the mode-locked laser after it has been frequency-shifted by the amount f_{AOM}. From the beat between the two combs an error signal used for suppression of the enviromentally-induced fibre noise is derived by phase comparison to a frequency reference. The noise suppression actuator is an adjustable delay line compensating for changes in the optical path length of the fibre link. MLL: mode-locked laser; SMF: single mode fibre; Det.: detection stage; DCF: dispersion compensating fibre; PC: polarization controller.

Figure 7.2 illustrates the optical phase detection principle. The Menlo Systems mode-locked fibre laser generates optical pulses at a repetition rate of $f_{\mathrm{r}} = 100$ MHz over an optical bandwidth of approximately 100 nm (this corresponds to approximately 10^5 optical modes). The frequency of each optical mode can be described by $f_m = mf_{\mathrm{r}} + f_0$ where m is an integer, f_r is the repetition rate and f_0 is the carrier-envelope offset frequency. Both f_{r} and f_0 are stabilized to a 10 MHz signal from a hydrogen maser. Approximately 10 m of SMF-28 fibre are used to connect the experimental setup to

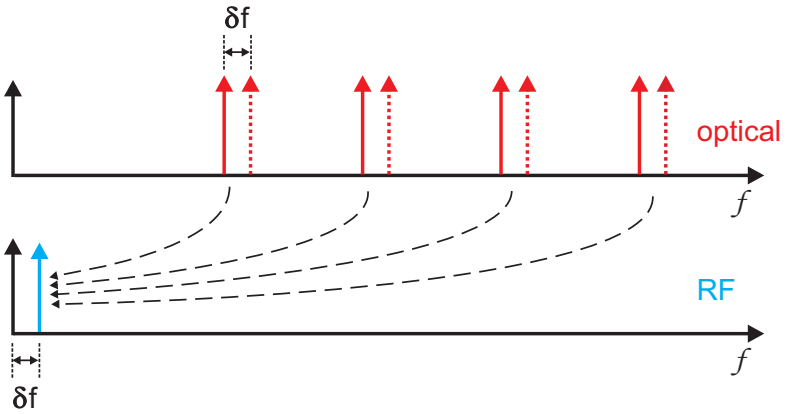

Figure 7.3: Illustration of the generation of the lowest frequency beat between the local and returned comb. Higher frequency beats are not shown as they are not used in the experiment described here.

the laser source, which is located in a different part of the laboratory, broadening the pulse duration to approximately 17 ps before it enters the first power splitter. The 90% output of this 90:10 splitter is used to propagate the comb over 7.7 km of spooled single mode fibre (SMF) to the receiver end where a portion is returned to the transmitter end via the same fibre. Both ends of the fiber are located in the same laboratory so that the accuracy and stability of the signal that has travelled 7.7 km can be compared to the signal injected at the input of the fiber. The forward and backward travelling pulse trains from the mode-locked laser are separated using optical circulators (CIR1, CIR2) and 200 m of dispersion compensating fiber recompress the pulses to a duration of less than 100 ps. A free-space delay line is adjusted to achieve appropriate temporal overlap between the local (10% output of the power splitter) and the returned pulse trains. The returned frequency comb is combined with the original comb after the latter has been frequency shifted by $f_{\mathrm{AOM}} = 104$ MHz using an acousto-optic modulator (AOM) and the beat notes between their optical modes are detected with a photodiode. The total length of out-of-loop fiber in the measurement setup was a few metres, with the lengths travelled by the shifted and unshifted combs being similar in order to improve the common-mode rejection of environmental effects that could increase the measurement noise floor.

The Doppler-induced frequency shift on the optical comb mode $f_p = pf_r + f_0$, due to environmental perturbations on the fiber, can be written as

$$\delta f_p = \frac{1}{c}\frac{d[n(f_p)l]}{dt} f_p \tag{7.2.1}$$

where $n(f_p)$ and $d[n(f_p)l]/dt$ are the refractive index of the fiber and the instantaneous rate of change of the optical path length at frequency f_p, and c is the speed of light in vacuum. This frequency shift can be measured by detecting any of the beat frequencies $f_{\mathrm{b}p,m}$ arising from interference between the optical modes f_m and f_p of the AOM-shifted

and unshifted combs respectively:

$$f_{\mathrm{b}p,m} = f_{\mathrm{AOM}} - (f_p - f_m) - \frac{1}{c}\frac{d[n(f_p)l]}{dt} f_p \tag{7.2.2}$$

where we have assumed that the optical frequencies generated by the mode-locked laser are constant over the round trip time so that no noise is detected due to the self-heterodyne effect. This assumption is reasonable when self-heterodyne noise is lower than the phase noise introduced by enviromental perturbations to the fiber link, which is the case in our experiment for frequencies within the feedback bandwidth of approximately 2 kHz. These detected beats give rise to a current from the photodetector that can be described as

$$i(t) \propto \sum_{m=m_1}^{m_2} \sum_{p=p_1}^{p_2} \cos 2\pi \left[f_{\mathrm{AOM}} - (p-m) f_r - \frac{1}{c}\frac{d[n(f_p)l]}{dt} f_p \right] t \tag{7.2.3}$$

where m_1, m_2, p_1 and p_2 define the range of comb modes that contribute to the signal. Eq. 7.2.3 illustrates the gain in sensitivity achieved using optical rather than microwave phase detection techniques; the frequency shift due to changes in the optical path length is greater by a factor $f_p/f_r \sim 10^6$. In our experiment f_{AOM} was 104 MHz, f_r was 100 MHz and we chose to detect the lowest frequency beat such that $p = m + 1$, in which case

$$i(t) \propto \sum_{p=p_1}^{p_2} \cos 2\pi \left[f_{\mathrm{AOM}} - f_r - \frac{1}{c}\frac{d[n(f_p)l]}{dt} f_p \right] t. \tag{7.2.4}$$

For simplicity, Eqs. 7.2.3 and 7.2.4 assume that all comb modes have the same power and interfere with the same phase offset. In reality, the residual dispersion of the fiber link means that destructive interference between the beat frequencies will reduce $i(t)$, reducing the signal-to-noise ratio of the extracted beat signal. However, this effect will be partially counteracted by variations in mode power across the spectrum, and in practice a sufficiently high signal is detected.

The lowest frequency beat is phase compared with a maser-referenced RF source, generating an error signal for the fibre noise cancellation. The error signal so constructed can then be used to drive optical actuators capable of counteracting the total optical path length changes of the optical link induced by acoustic noise and temperature changes over time. If sufficient gain is available around the feedback loop, the enviromentally-induced fibre noise can be efficiently cancelled.

7.3 Experimental test of the detection principle

In order to test the sensitivity of the optical detection technique, I have run an experiment with the interferometric setup depicted in figure 7.4. The output of the frequency comb is split into two measurement arms with a power splitter. One output of the power splitter is connected to an AOM driven with an RF signal of frequency 97 MHz. A fibre stretcher is connected to the other output of the power splitter and is followed by a free space delay line and a variable attenuator used to balance the optical power in the two arms. The two measurement arms are recombined using a 50:50 power combiner. The beat δf between the combined (shifted and unshifted) combs is detected with a photodetector. The lowest frequency beat, 3 MHz in this experiment ($f_b = f_r - f_{\mathrm{aom}}$=100 − 97 = 3 MHz), is selected with a 5 MHz low pass filter and phase compared to a frequency synthesizer which is referenced to the same reference as the AOM driver. The voltage on the fibre stretcher is modulated with a 4 Hz sinusoid. The phase fluctuations on the 3 MHz signal resulting from the 4 Hz modulation of the optical path length changes are measured by recording the output voltage of the mixer with a digital scope. The phase cycles as a function of the fibre stretcher modulation are shown in figure 7.5. By counting the number of phase cycles for a given change in the optical path length it is possible to measure the sensitivity of the optical phase detection technique. The frequency of the modulation was chosen to be faster than the temperature-induced optical path length imbalance between the the two arms of the interferometer whilst, at the same time, inducing a sufficiently low number of phase cycles to be easily counted on a scope. The optical path length displacement of the fibre stretcher, as shown in its datasheet, is 8.1 μm/V at 1.55 μm. Driving the fibre stretcher with a triangular wave of amplitude 4 V we should expect a phath length change of

$$\delta l = 8.1\,\frac{\mu\mathrm{m}}{\mathrm{V}} \cdot 4\,\mathrm{V} = 32.4\,\mu\mathrm{m} \tag{7.3.1}$$

which corresponds to 20.8 cycles at the centre wavelength of the frequency comb 1.56μm. This is in reasonable agreement with the observed 21.5 cycles shown in figure 7.5. The discrepancy can be attributed to the asymmetry of the optical spectrum around the supposed center wavelength (1.56 μm in this case) and the approximate value of the optical path legth displacement stated in the datasheet. The Doppler-induced frequency shift due to the linear change in phase in 250 ms will be

$$\delta f = \frac{21.5\,\mathrm{cycles}}{0.25\,\mathrm{s}} = 86\,\mathrm{Hz} \tag{7.3.2}$$

From this experiment it can be deduced that the comb interferometer can be thought of as a single optical carrier interferometer where the wavelength of the optical carrier is the average wavelength of the optical frequency comb.

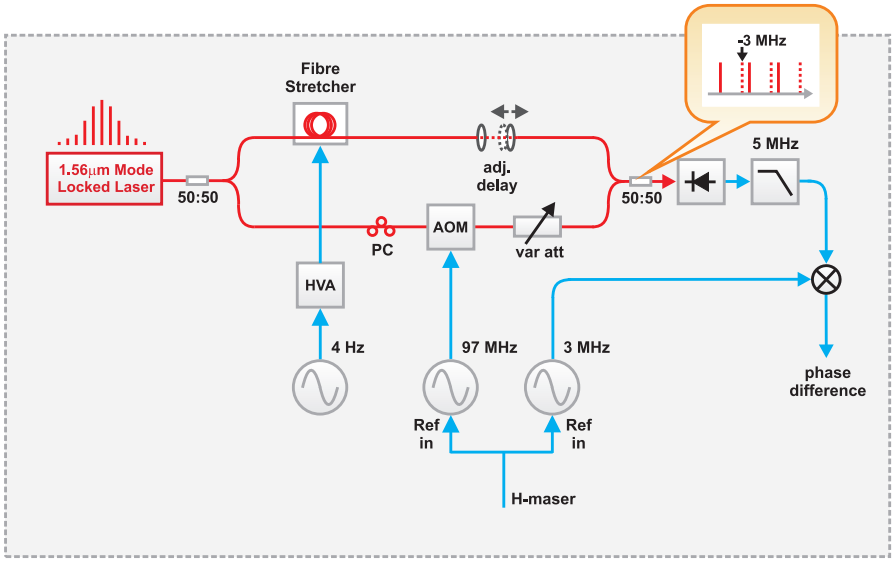

Figure 7.4: Experimental setup for the proof of the optical phase detection principle. HVA: high voltage amplifier; Var. att.: variable attenuator; PC: polarization controller. The fibre stretcher imposes optical path length changes which are detected as phase changes of the beat between the shifted and unshifter combs.

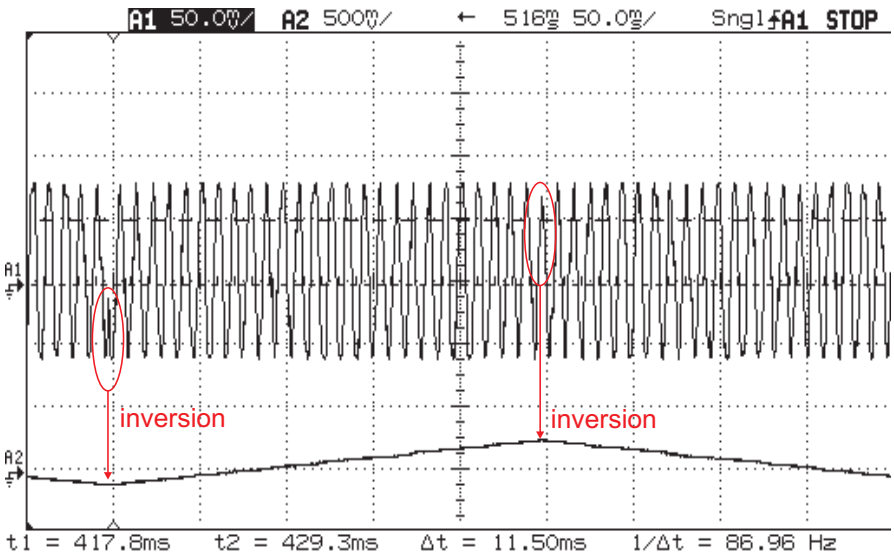

Figure 7.5: Time domain data from a digital oscilloscope showing the variations in the phase of the detected beat between the shifted and unshifted comb as a function of the driving voltage of the fibre stretcher.

7.4 Comb transfer experimental setup

Figure 7.6 shows the experimental setup for the transfer of an optical frequency comb over a fibre 7.7 km long. The mode-locked laser used as the optical source for this experiment is the same as that used for the fibre transfer experiment described in Chapter

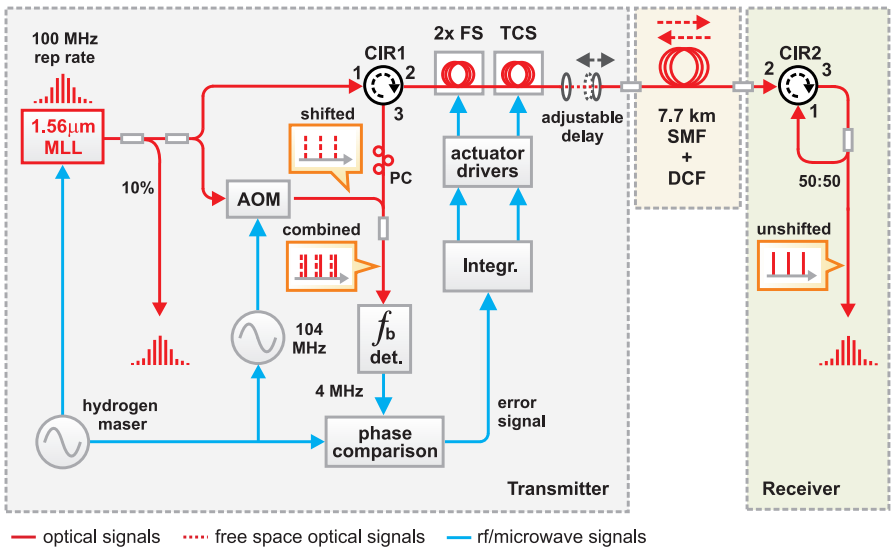

Figure 7.6: Block diagram of the experimental set-up for the transfer of an optical frequency comb over 7.7 km of spooled fibre. Each repetition rate detection module consists of a photodiode followed by a tunable band pass filter and cascaded RF amplifiers. The block marked as f or ϕ comparison is either a phase detector followed by a fast Fourier transform analyser or a single sideband modulator-based frequency counting system. MLL: mode-locked laser; CIR: optical circulator; SMF: single mode fibre; Int.: integrator; Det.: detection stage; DCF: dispersion compensating fibre; EDFA: erbium-doped fibre amplifier; PC: polarization controller.

5. Both f_r (100 MHz) and f_0 are stabilized to low noise synthesizers referenced to the H-maser. It is interesting to note that the stabilization of f_0 is not strictly required for the noise cancellation system to work as changes in its value are common mode to the reference (local comb) and measurement (returned comb) arms for time scales much longer than the delay imposed by fibre. However, the carrier envelope offset frequency was stabilized during the fibre transfer experiment as this was required for other experiments taking place at the same time involving the same optical frequency comb. Short tests of the transfer stability when f_0 was not stabilized have comfirmed the high level of common mode rejection between the reference and measurement arms.

A 10 m fibre is used to link the comb to the optical table where the optical setup is mounted, broadening the sub-150 fs pulses available at the output of the mode-locked laser to a few tens of ps before they enter the experimental setup. A 50:50 splitter is used to generate the output for the reference arm. The pulse train from the other output of the power splitter is propagated through two fibre stretchers and a thermally controlled fibre spool, a free-space optical delay line, 7.7 km of single mode fibre and a dispersion compensating module matched to the length of the single mode fibre. At the receiver end, a circulator allows for half of the pulse train optical power to be delivered to a "user" and the other half to be re-injected into the fibre link. The returned frequency comb at the transmitter end is combined with that frequency-shifted by the AOM in

the reference arm. The AOM is driven with a frequency of 104 MHz and an RF power of approximately 25 dBm. As the repetition rate is 100 MHz, it follows that the lowest beat note between the local and returned frequency combs is at 4 MHz. This beat note is extracted by a detection stage consisting of an RF photodiode followed by a 5 MHz low pass filter and three cascaded amplifiers used to increase the detected RF power to a suitable level for the further processing stages.

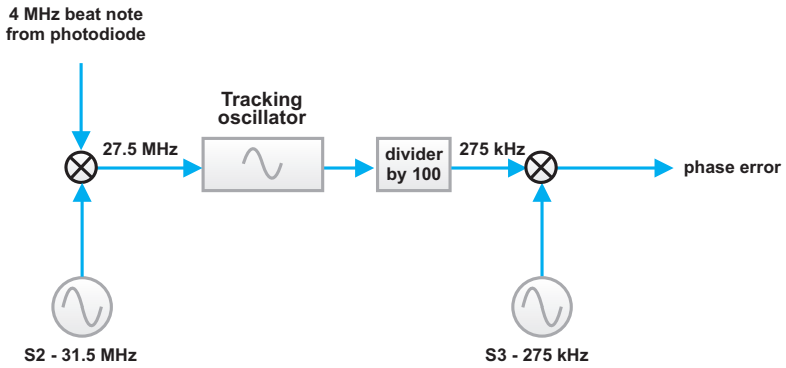

Figure 7.7: Block diagram of the frequency chain required to upconvert the 4 MHz beat to the acceptable frequency range of the tracking oscillator (20-30 MHz) employed to reject the broadband phase noise. The divider by 100 is used to reduce the detected phase fluctuations to the linear range of an analog phase detector (frequency mixer with inputs in quadrature)

A small frequency chain was required to convert the 4 MHz beat note up to the operating frequency range of a tracking oscillator[1] (25 to 35 MHz) and to reduce the phase fluctuations to within the linear phase range of an analog RF phase detector (mixer with inputs in quadrature). This frequency chain setup is shown in figure 7.7. The 4 MHz beat note is mixed with a synthsizer set at 31.5 MHz, generating two frequencies, 27.5 MHz and 35.5 MHz. A tracking oscillator is phase locked to the 27.5 MHz beat note. The tracking oscillator is followed by a digital frequency divider by 100 whose 275 kHz output is phase compared with a synthesizer set at the same frequency. The tracking oscillator was used to remove the broadband noise around the 4 MHz beat note in order for the digital frequency divider to operate correctly. The signal-to-noise ratio of the detected beat note was in fact only of the order of 40 dB over a 10 kHz measurement bandwidth. Although such a low signal-to-noise ratio is more than sufficent for efficient fibre noise cancellation, the broadband noise around the beat note renders the zero-crossing detection process, required for the correct functioning of the digital divider, difficult. The tracking oscillator greatly reduced the noise at offset frequencies higher than the 200 kHz locking bandwidth whilst preserving the measured noise introduced by the fibre. The digital divider was required in order for the large detected phase fluctuations to be reduced to within the linear range of an analog phase detector. The

[1] A tracking oscillator is esentially a voltage controlled oscillator phase locked to an input signal. This is usually called phase-locked loop (PLL) in the field of RF and microwave electronics, but for some historical reason it is called a tracking oscillator in the laser community.

division ratio was chosen to be 100 leading to the phase comparison being performed at 275 kHz.

The error signal available at the output of the analog phase detector is then fed to two integrators driving the optical actuators required for the fibre noise cancellation. The two fibre stretchers are used to compensate for fast phase fluctuations (up to a few kHz) arising primarily from acoustic noise, whilst a thermally controlled spool compensates for slow fluctuations due to thermal effects in the fibre link. The fibre stretchers are driven by high voltage amplifiers and the fibre spool is driven by a power amplifier (2 A). The fibre stretchers and the thermally controlled fibre spool are the same as used in the fibre transfer experiment over the Aurora network, described in chapter 6.

7.5 Comb mode stability measurement principle

Whilst the measurement of the repetition rate can be performed in a relatively straightforward way using a simple photodetector, as was done for the previous microwave transfer experiments, the measurement of the optical mode frequency requires a more sophisticated approach which is illustrated in figure 7.8. A continuous wave (CW) laser produces an optical carrier at a frequency f_1 which is combined with the optical frequency comb before and after the fibre link. At each end of the fibre, a number of frequencies are generated with a photodetector as a result of the beating between the CW carrier f_{cw} and the optical modes. The lowest of these beat frequencies is that arising from the interference of the CW carrier with its nearest comb mode. At the local end of the fibre we have

$$f_{\mathrm{loc}} = f_m - f_{\mathrm{cw}} \tag{7.5.1}$$

and at the remote end

$$f_{\mathrm{rem}} = f_m + \delta f - f_{\mathrm{cw}} \tag{7.5.2}$$

where f_m is the frequency of the m optical mode and δf is the Doppler-induced frequency shift due to optical path length changes in the fibre. In the absence of fibre noise, the two beat frequencies are identical i.e.

$$f_{\mathrm{rem}} = f_{\mathrm{loc}} \tag{7.5.3}$$

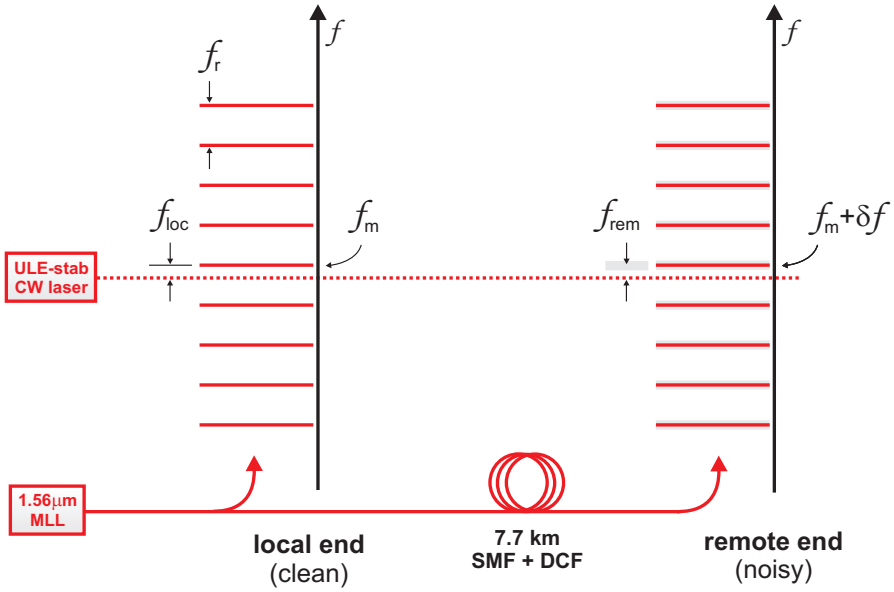

Figure 7.8: Operating principle of the comb mode stability measurement. The same comb mode is compared to a stable laser before and after propagating into the fibre.

The noise introduced by the fibre will make the local and remote beat frequencies differ by the amount δf. By measuring this difference the frequency instability of the optical modes caused by the fibre can be measured.

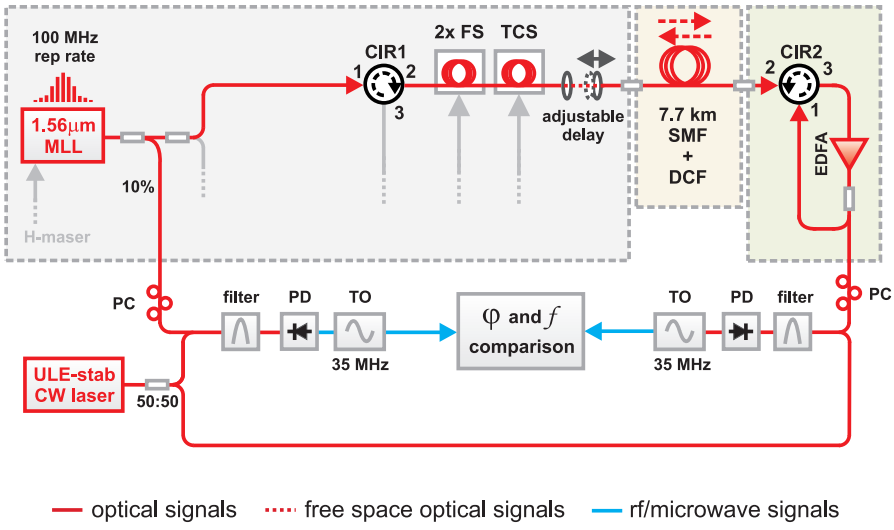

Figure 7.9: Block diagram of the experimental setup for the measurement of the stability and accuracy of the optical comb modes delivered at the remote end of the fibre. MLL: mode-locked laser; CIR: circulator; FS: fibre stretcher; TCS: thermally-controlled spool; SMF: single mode fibre; DCF: dispersion compensating fibre; PD: photodiode; EDFA: erbium-doped fibre amplifier; TO: tracking oscillator; ULE-stab CW: ultralow expansion cavity-stabilized continuous wave laser.

7.6 Measurement of the stability and phase noise of the optical modes

At each end of the fibre, a 35 MHz note is generated from the beat between the comb and a CW laser (f_{loc} and f_{rem}) as described in the previous section and shown in Figure 7.9. At the remote end, an optical amplifier was used in order to achieve a suitably high signal-to-noise ratio of the beat. Approximately 4 mW of average power is launched into the fibre and, without amplification, only 180 μW are received at the remote end, corresponding to a power of the order of 1 nW per mode. Tracking oscillators were used on the beats in order to correctly drive the digital phase detector and the frequency counters. The feedback bandwidth of these tracking oscillators was in excess of 200 kHz. From the comparison of these two beats both the phase noise and frequency stablity can be measured.

7.7 Frequency stability of the optical modes

The transfer stability of the optical modes was measured by comparing f_{loc} and f_{rem} in two ways:

1. by phase comparison using a digital phase detector;
2. by synchronously counting their frequency.

These two methods provide stability measurements over two different measurement bandwidths. When the digital phase detector is used, the stability is calculated by conversion of its output voltage, acquired using a digital voltmeter. The digital voltmeter has a 7 Hz low pass input filter, so phase fluctuations much faster that this frequency (or in other words, the phase noise at offset frequencies from the carrier much higher than 7 Hz) will not be measured. The filter therefore sets the measurement bandwidth. When the frequency counters are used, the measurement bandwidth is that of the tracking oscillator (200 kHz) as faster phase fluctuations will not be tracked. The frequency stability results will therefore differ when the frequency counters are used as more noise is integrated compared to the case where the digital phase detector is used. For long integration times, when the fast noise processes are averaged out, the stability should in principle reach the same value for the two measurement techniques (and I will show that this is indeed the case). In the case of 200 kHz measurement bandwidth, the measured noise contains not only the residual noise of the fibre, but also the noise of the laser source detected by the self-heterodyne effect arising from the delay introduced by the fibre. This noise can be observed as dips and bumps in the phase noise of Figure 7.12 in the 10 to 100 kHz offset frequency region. The frequency stability measured with

a 200 kHz measurement bandwidth therefore represents the worst case for the actual transfer stability.

The stability measured with a 7 Hz measurement bandwidth allows the stability that would be available at the user end if the comb was stabilized to an optical reference to be estimated. In the experiment described here, the comb is stabilized to a hydrogen maser and consequently the optical modes jitter over a few hudreds kHz, orders of magnitude more than the measurement bandwidth. However, the residual noise of the fibre is measured as the difference of two 35 MHz beat notes detected at the local and remote end of the fibre which are highly correlated for offset frequencies up to 7 Hz from the carrier. At much higher offset frequencies their noise becomes decorrelated because of the delay introduced by the fibre. As the noise for offset frequencies up to 7 Hz is dominated by the measurement setup and not by the self-heterodyne detection of the laser noise, it would be found to be the same if an optical reference was used instead of the hydrogen maser to stabilize the comb (and thus the jitter was reduced to less than the 7 Hz measurement bandwidth).

The fractional frequency stability of a sample optical mode delivered to a "user" after 7.7 km of fibre is shown in figure 7.10. The frequency stability calculated from the phase data is 4×10^{-17} at 1 s and approximately 2×10^{-18} for timescales of a few thousand seconds. When the frequency counters are used, so that the measurement bandwidth is 200 kHz, the frequency stability is 5×10^{-15} at 1 s and reaches a few parts in 10^{18} at a few thousand seconds.

The measurement noise floor was measured by replacing the fibre spool (SMF and DCF module) with an optical attenuator set at an equivalent loss. As shown in figure 7.10, the transfer stability almost exactly matches the measurement noise floor. This means that, for a 7 Hz measurement bandwidth, all the phase noise within the measurement bandwidth is cancelled by the fibre noise suppression loop. The measurement noise floor is limited by the temperature imbalance between the two arms of the phase detection setup where the pulse trains travel through non-common paths. This issue is illustrated in figure 7.11. The phase detection setup consist of two couplers, an AOM and a circulator. Each optical component is provided by the manufacturer with at least 1 m of fibre for each input/output port. The non-common path travelled by the pulse trains is therefore a few metre long and the optical path length changes due to thermal effects in this fibre limit the minimum phase difference that can be measured. Further improvements of the measurement noise floor could be achieved by higher integration of the phase detection setup, as will be discussed later.

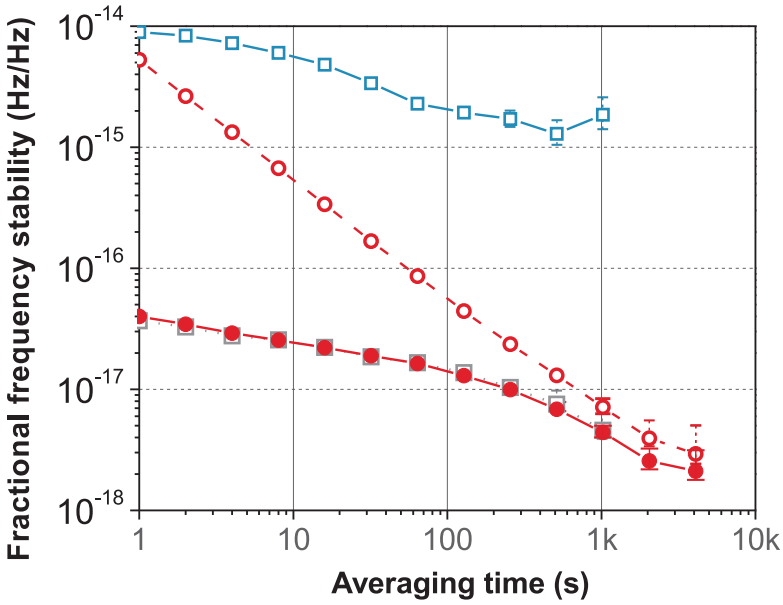

Figure 7.10: Fractional frequency stability of a sample optical mode delivered at the user end after 7.7 km of fibre. The results are shown for 7 Hz (solid red circles) and 200 kHz (void red circles) measurement bandwidth. Solid grey circles: measurement noise floor; Empty blue squares: phase noise cancellation off.

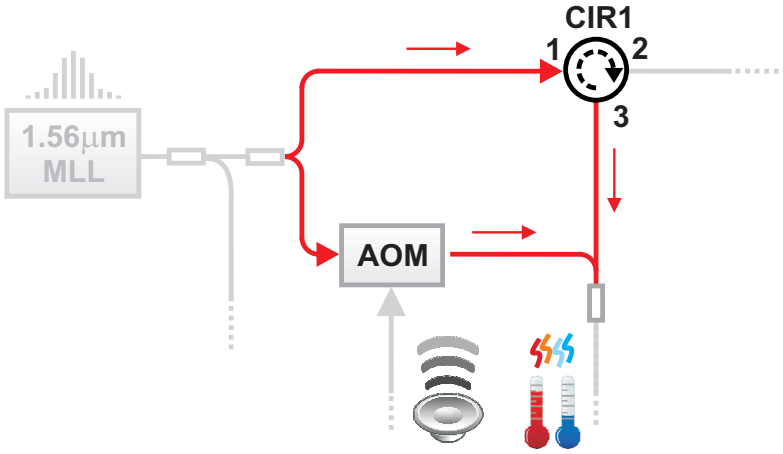

Figure 7.11: The temperature imbalance between the two non-common paths of the optical interferometer used to detect the fibre phase noise can limit the measureable stability.

7.8 Phase noise of the optical modes

The phase noise of a sample comb mode delivered at the "user" end is shown in figure 7.12. The noise introduced by the fibre is reduced by up to 55 dB when the phase noise suppression loop is activated and is −40 dBc/Hz at 1 Hz offset from the carrier. In

the phase noise plot the dips due to the self-heterodyne effect are visible. These fall at multiples of the inverse of the delay τ introduced by the fibre link. As the total length of the fibre is 8.7 km it follows that

$$\tau = \frac{8.7 \cdot 10^3\,\mathrm{m}}{2 \cdot 10^8\,\mathrm{m/s}} = 43.5\,\mu\mathrm{s} \tag{7.8.1}$$

therefore the self-heterodyne minima in the phase noise will be at multiple of the frequency f_{sh}

$$f_{sh} = \frac{1}{43.5 \cdot 10^{-6}\,\mathrm{s}} \simeq 23\,\mathrm{kHz} \tag{7.8.2}$$

where the speed of light in the fibre has been approximated to 2/3 of that in vacuum. The calculated value is in very good agreement with the measured data. The grey trace of figure 7.12 shows the limit to the transfer phase noise performance imposed by the delay-unsuppresed fibre noise as per (Williams et al., 2008). In figure 7.13 the evolution of the phase fluctuations over time is shown and they are measured to be less than 100 rad over 35000 s, corresponding to appoximately 16 optical cycles (corresponding to less than 5×10^{-4} cycles over 1 s).

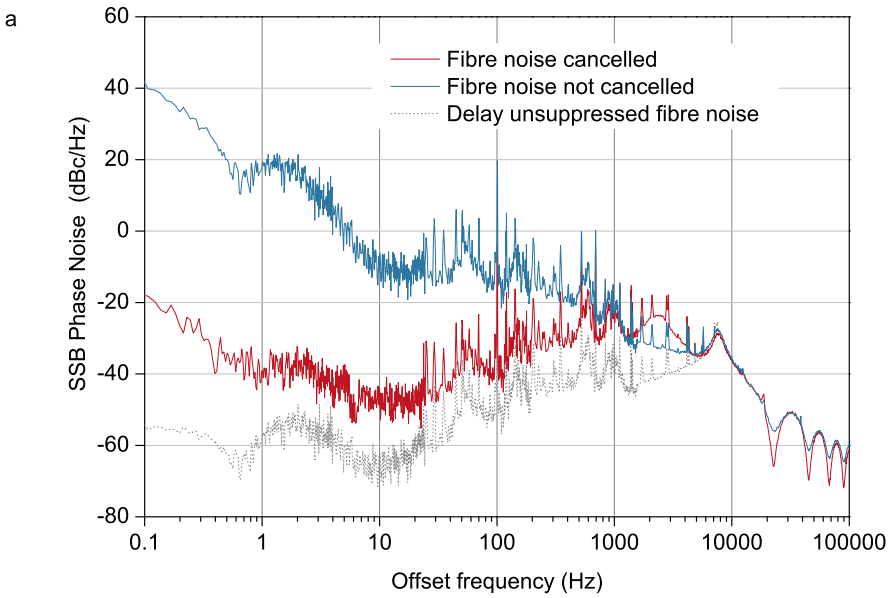

Figure 7.12: Single sideband phase noise of a sample optical mode delivered at the user end after 7.7 km of fibre when the fibre noise cancellation loop is off (blue trace) and off (red trace). The grey trace is the delay-unsuppressed fibre noise limit.

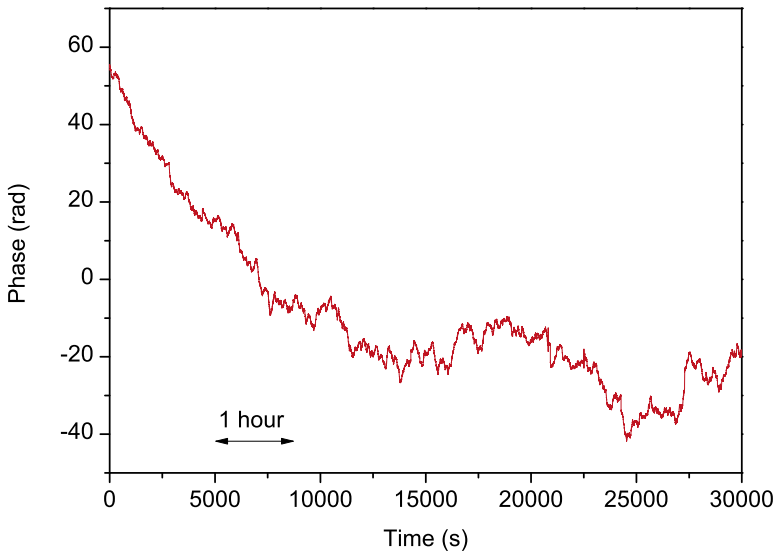

Figure 7.13: Evolution of the phase of a sample optical mode at the user end over 35,000 s (9.7 hours) with the fibre noise cancellation loop active. The maximum phase deviation over this time is less than 100 rad.

7.9 Stability of the CW laser

It is important to note that the stability of the CW laser is, to first approximation, irrelevant for the measurement of the comb mode stability and accuracy as it is common-mode between the two measurement arms. Second order effects could potentially arise if the CW laser instability were high emough to reveal any frequency-dependent behaviour of the optical and microwave components. For this experiment, a cavity stabilized laser whose frequency would drift by no more than 1.5 Hz over a second was available, rendering any second-order effects negligible.

7.10 Measurement of the stability and phase noise of the optical mode spacing

The measurement of the both the stability and phase noise of the mode spacing is performed by phase comparing the 80^{th} harmonic (8 GHz) of f_{r} at the receiver end of the fibre with that detected directly at the output of the laser as shown in figure 7.14. The phase comparison is achieved using a microwave mixer with the input signals in quadrature and the output voltage of which is measured using a digital voltmeter with a measurement bandwidth of 7 Hz. In this case the optical power launched into the fibre was approximately 10 mW and no optical amplifier was required at the remote end as over 130,000 optical modes contribute to generate the same 8 GHz signal (without the optical amplifier the comb's optical bandwidth spans approximately 100 nm).

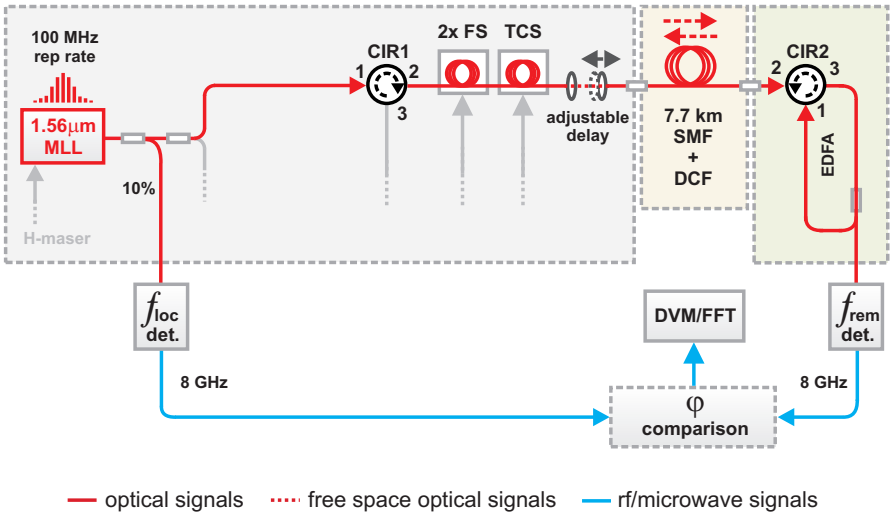

Figure 7.14: Experimental setup for the measurement of the stability and phase noise of the transferred 80^{th} harmonic of the repetition rate over 7.7 km of optical fibre.

7.11 Frequency stability of the optical mode spacing

The transfer frequency stability is shown in figure 7.15 and is calculated by converting the phase data into frequency and then applying the Allan deviation formula. Over a measurement period of two hours the Allan deviation is 1.8×10^{-15} at 1 s and 7×10^{-17} at 100 s. Over much longer time scales, the overall stability degrades slightly to 2.2×10^{-15} at 1 s and 1×10^{-16} at 100 s. The higher instability measured when longer data sets are taken is due to degradation of the signal-to-noise ratio of the error signal caused by thermally-induced polarization changes in the fibre. This problem could in future be overcome using either an optical amplifier to increase the signal to noise ratio or introducing an automatic control system to stabilize the polarization of the returned pulse train. The degradation is however modest and, depending on the application, it might be considered acceptable when traded off against the extra complexity required to implement the polarization control.

7.12 Phase noise of the optical mode spacing

The measured phase noise is shown in figure 7.16. When the noise cancellation is activated, the measured phase noise is −91 dBc/Hz at 1 Hz offset from the carrier, very close to that measured when the 7.7 km SMF and the DCF are replaced by an attenuator set to provide the same overall loss. The noise floor was measured by replacing the fibre spools with a variable attenuator set to an equivalent loss.

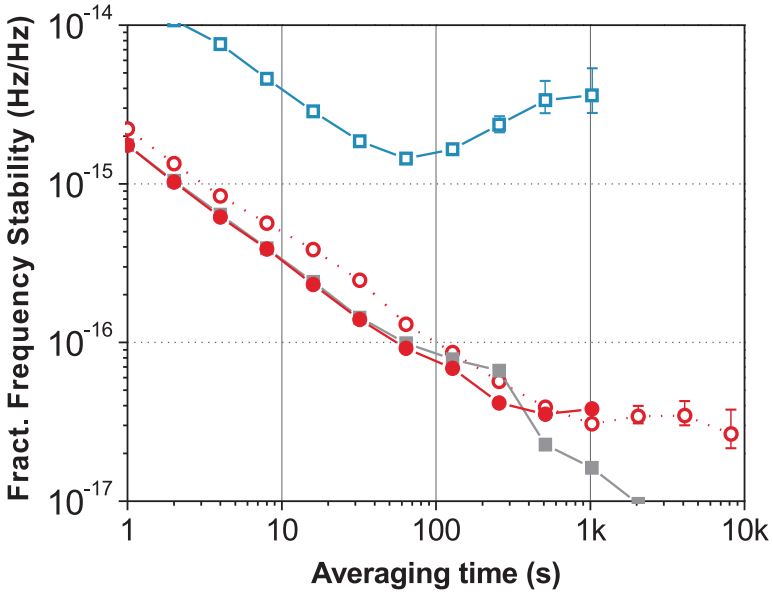

Figure 7.15: Fractional frequency stability of the mode spacing (repetition rate) delivered at the user end after 7.7 km of fibre measured on the 80^{th} harmonic (8 GHz) when the fibre noise cancellation loop is off (void blue squares) and on (filled red circles). The void red circles show the stability degradation due to a reduction of the signal-to-noise ratio of the detected beat (see text). Solid grey circles: measurement noise floor.

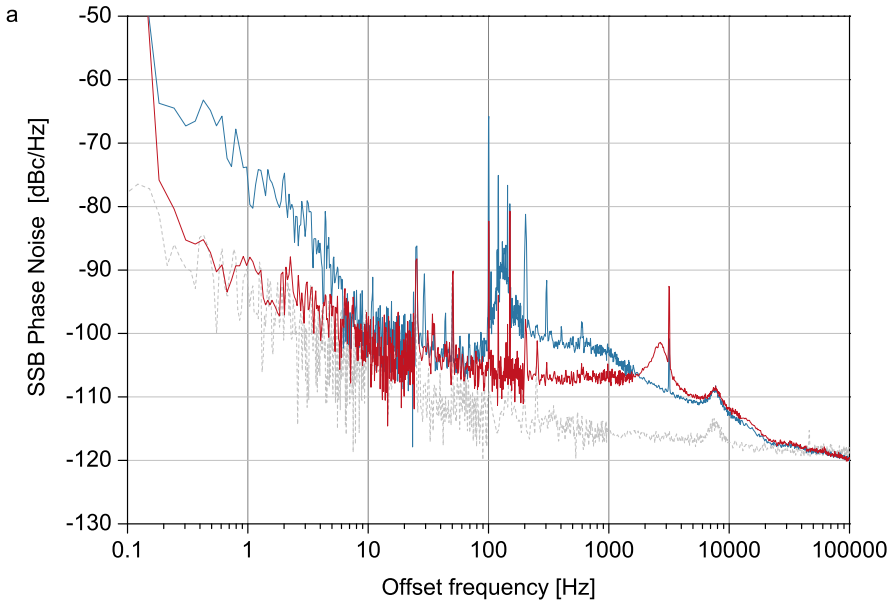

Figure 7.16: Phase noise of the 8 GHz harmonic delievered at the used end.

7.13 Effect of air circulation

The temporal overlap between the local and returned pulse train required is achieved with a free space optical delay line. As the repetition rate is 100 MHz, the pulses repeat every 10 ns which corresponds to a travel distance of approximately 2 m. Very coarse adjustment of the delay was achieved by joining together short FC/APC connectorized patch cables. A finer adjustment was achieved with a home built delay line. I constructed this by mounting two fibre coupled collimator lenses on linear translation stages with micrometer adjustments which in turn were mounted on clamps that could be slid on an aluminium profile beam. Using the sliding clamps the delay could be changed by up to 50 cm with a precision of approximately 1 mm, whilst the translation stage allowed sub-mm adjustments. The loss was between 2 and 4 dB depending on the distance between the collimators, with the highest loss measured for the longest distance due to the divergence of the beam. Using short optical patch cords (10 to 30 cm) before the delay line the distance between the collimator in the free space delay line was generally kept below 30 cm, mantaining the transmission loss to less than 3 dB. The initial temporal delay between the pulse trains (and conversely the correction that needs to be applied to overlap them) can be estimated easily by looking at the spectrum of the unfiltered output of the photodiode used for the 4 MHz beat detection. The Fourier transform of two interlaced pulse trains is still a frequency comb but intensity modulated by a sinc function with minima at multiples of the inverse of the delay between the pulses. The frequency spectrum obtained for different delays is shown in figure 7.17.

The delay is adjusted coarsly until the first minima is no longer observable within the bandwidth of the photodetector (a few GHz) and the 4 MHz beat note starts to appear. The fine adjustment is achieved by looking at the intensity of the detected 4 MHz beat note which is maximum when the pulses overlap best in time. Once the fibre noise cancellation loop is activated, the pulse overlap condition is automatically preserved as the optical path length is stabilized.

The free space delay line is intrinsically a potential source of instability. Vibrations and thermal expansions of the mounts and the lenses can cause changes in the alignment converting into power fluctuations of the propagated pulse train. Also, dust particles carried by turbulent air can modulate the beam intensity. The mechanical stability of the delay line has shown not to be a limiting factor at the stability level investigated here. The turbulence of the air has instead proven to be a crucial factor in limiting the achievable transfer stability. Simply reducing the air circulation across the beam with a protective tube was enough to improve the amplitude noise of the extracted 8 GHz harmonic of the repetition rate between 0.1 and 20 Hz by up to 13 dB. The amplitude noise was measured using a microwave power detector. The results of this test are shown in figure 7.18.

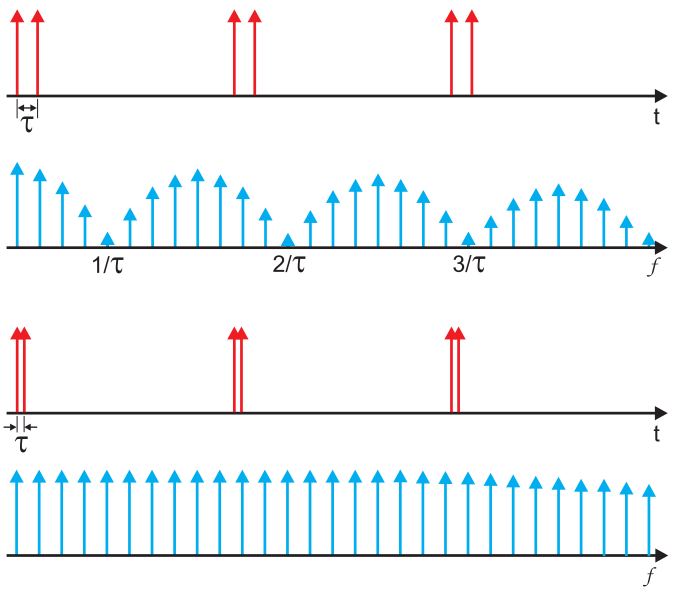

Figure 7.17: Change in the envelope of the microwave spectrum of two interleaved pulse trains detected with a photodetector. Spectral minima appear at a multiple of the inverse of the delay between the first and the second pulse train. The delay between the optical pulses can thus be estimated from the microwave spectrum by looking at the position of the minima. A quasi-flat envelope indicates that the pulses almost overlap in time.

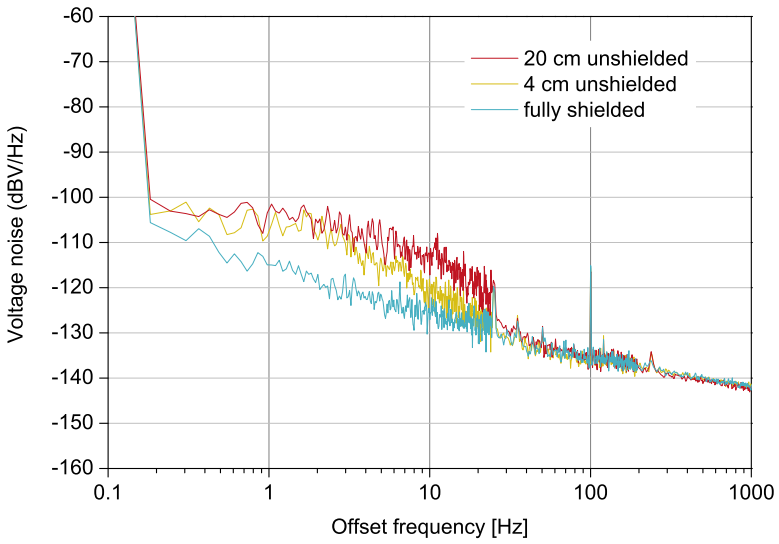

Figure 7.18: Amplitude noise of the detected 8 GHz harmonic for different levels of shielding of the free-space optical delay line. Leaving even a small fraction of the optical path length unshielded was enough to see a substantial rise in the AM noise between 0.1 and 20 Hz.

7.14 Performance comparison with the single carrier transfer technique

Figure 7.19 shows the fractional frequency stability of the optical modes delivered at the remote end of the 7.7 km link in comparison with that achieved by other groups (LPL-SYRTE, PTB, NIST, NMIJ) when a single optical carrier is transmitted over links of different length. Comparing results achieved over different lengths of fibre and different enviromental conditions is not straightforward as the noise profile can be rather different, but it is however interesting to get a better interpretation of the physical meaning of the results. The measurement of the transfer stability using the 7 Hz measurement bandwidth highlights that, if the comb was stabilized to an optical reference, tens of thousands of optical modes could be transferred with a fractional stability comparable to that achieved when only one optical carrier is transferred.

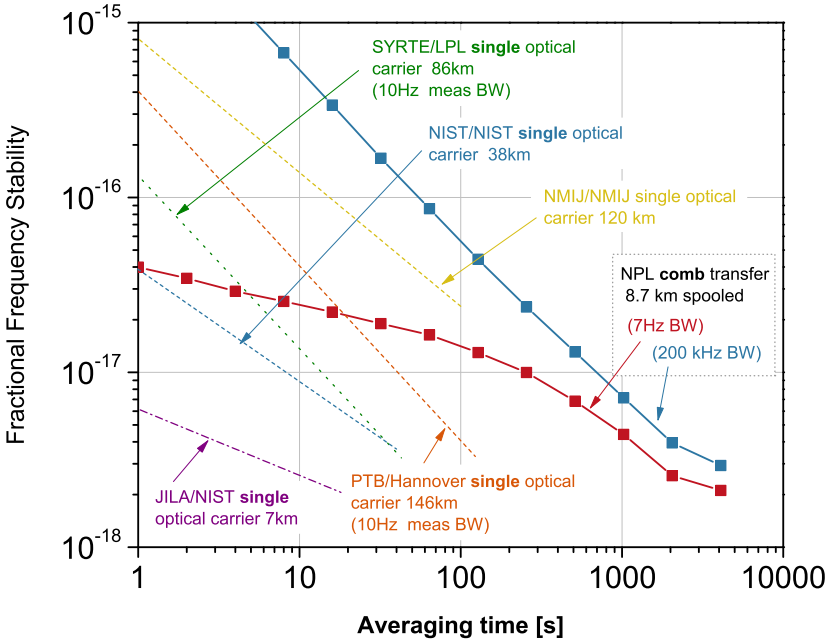

Figure 7.19: Performance comparison between the optical comb transfer and the single carrier transfer techniques. The results shown in this chapter demonstrate that it is possible to transfer tens of thousands of coherent modes at a comparable level of stability to that achieved when only one optical carrier is transferred.

The experiment that has been described in this chapter is, to my knowledge, the first demonstration that a frequency comb can be disseminated to a remote location with a stability and accuracy comparable to that of the best optical references. A frequency comb extending over an optical bandwidth of approximately 30 nm, corresponding to over 35,000 optical modes, was transferred over 7.7 km of fibre whilst preserving the

frequency stability and accuracy of its optical modes to the level of parts in 10^{-18} and the mode spacing to parts in 10^{-17}. The results enable a range of possible applications, which are described in the next chapter.

Chapter 8

Conclusions

In this thesis I have shown results of three frequency transfer experiments over optical fibre. In the first and the second experiment (Chapters 5 and 6 respectively), the regular spacing of the teeth of an optical frequency comb was used to transfer a microwave reference over optical fibre links of many tens of kilometers. The link consisted of 50 km of spooled fibre in the first experiment and 86 km of installed fibre link in the second. The experiments demonstrated that the noise of the fibre can be suppressed to a level such that the frequency spacing of the optical comb delivered at the user can be mantained at a stability of parts in 10^{-15} at 1 s and up to 4×10^{-17} at longer time scales. Building on these experiments, I have then looked at another crucial issue for dissemination of optical frequency combs over fibre. Whilst the first two experiments demonstrated that the frequency spacing between the teeth of the transferred comb could be preserved to a very high level of stability and accuracy, this could potentially not be true for the individual comb frequencies as the whole comb could potentially "wander" in frequency. The third experiment showed that the whole comb structure could indeed also be preserved, when transferred over a many-km optical link, to a very high level of accuracy and stability which is compatible with the best frequency references available today. The third experiment also applied a different detection technique for the measurement of the noise of the fibre, which was measured in the optical domain rather than after optical-to-microwave conversion, providing substantially higher sensitivity.

The main outcome of all three experiments is that optical frequency combs can be disseminated over long lengths of optical fibre whilst preserving their structure at a very high level. In the following paragraphs I discuss possible applications of these frequency transfer techniques as well as outline future work.

8.1 Dissemination within research laboratories

Within metrological laboratories such as NPL, accurate and stable combs could be disseminated from a "master" optical frequency comb referenced to an optical clock, or an optical reference such as an ULE cavity-stabilized laser, disseminating stable optical frequencies to a number of other labs. The master comb could reside in a quiet environment together with the optical source, such that continuous operation and consistent perfomance can be guaranteed. At each end of the fibres, the users will be able to lock their lasers to stable comb lines (therefore sub-Hz linewidth sources can be generated in each lab without the need for expensive ultra stable cavities), or to use the comb as their high precision optical frequency counter or to make use of the delivered pulse train for timing applications. In a way, such distribution of referenced optical frequencies can be seen as the optical equivalent to the standard 10 MHz reference distribution available in most metrology labs. Within the same laboratory, the typical distances over which the comb would need to be disseminated would probably be no longer than 100-300 m. Over these distances no optical amplification would be required so that wide combs could be transferred. Using standard SMF-28 fibre, over 400 nm of comb could be transferred (from 1200 to 1600 nm) with only moderately low modulation of the intensity over the entire spectrum due to the fibre's wavelength-dependent loss. Frequency combs in the visible part of the spectrum could also be distributed with appropriate fibre.

The dissemination of a comb could actually provide not only an optical reference, but could also act a replacement of the standard 10 MHz reference. The 10 MHz signal used in metrology labs is usually generated by good references, which in the best case is a hydrogen maser. Whilst the long term stability of the maser approaches 1×10^{-15} (at 1 day), the short term stability is much poorer, normally $2 \times 10^{-13}\,\tau^{-1/2}$. For some applications this is insufficient (as in the case of caesium atomic fountains). An ULE cavity-stabilized laser provides a much higher level of stability which is generally below a few parts in 10^{15} between 0.1 and 100 s and, if the ultra-stable laser is part of an optical clock, the stability can be maintained for much longer time scales. If not only the stability but also the accuracy of the optical or microwave signals are a concern to the user, the optical reference will necessarily need to be referenced to an atomic transition.

8.2 Dissemination to remote research laboratories and industry

When a research fibre network is available, such as the JANET-Aurora or the Thames Valley Network (UK) or the cross-border dark fibre links between Austria, the Czech

republic and Slovakia (ACOnet, CESNET2 and SANET networks), the stable and accurate optical frequencies could be distributed to remote laboratories and industries, as illustrated in Figure 8.1. In this case the distances over which the comb need to be propagated will be much longer than in the case of the distribution between neighbouring laboratories. Optical amplifiers are therefore likely to be required, restricting the optical bandwidth that can be transmitted. In the case of industry, the bandwidth restriction will probably not be an issue as their optical region of interest would most likely match that available from optical amplifiers. In the case of research laboratories, if other regions of the spectrum are of interest, spectral broadening might be required and how this can be achieved is discussed later in this chapter. The small extra complexity could be well worthwhile if the user laboratory has no other means of generating accurate optical frequencies (which is usually the case for non-metrology laboratories).

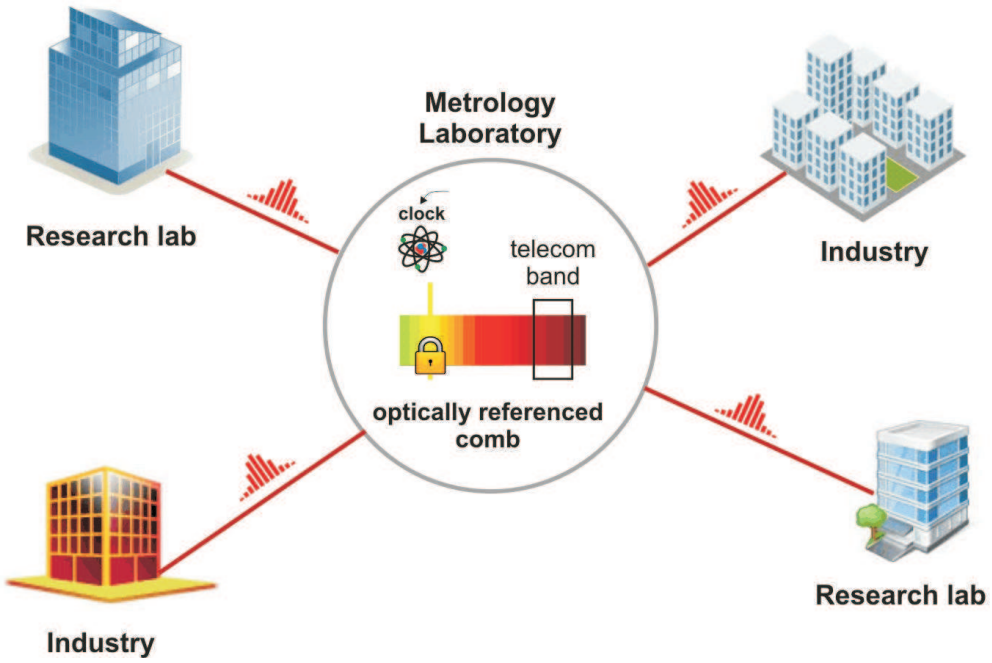

Figure 8.1: Possible future: dissemination of ultra stable optical frequencies to other research laboratories and industries.

8.2.1 Frequency comb transfer between satellites

This is a possible future application which I find particularly interesting as it exploits the full potential of the comb transfer technique. Scientists around the world are dreaming of putting optical clocks in space as this could enable a number of new experiments in the fields of fundamental physics, geodesy, astronomy, Earth observation, communications and global navigation, deep space ranging and communication and comparison of high accuracy ground-based optical clocks. The European Space Agency commissioned a feasibility study on the topic a few years ago and more discussions are likely to take

place in the near future. If optical clocks were deployed in space, this means that one or more optical frequency combs will be present in each satellite. With very limited extra complexity, the frequency comb could be transferred between the satellites, allowing simultaneous

1. microwave and optical frequency transfer;
2. time transfer;
3. ultra accurate distance measurements.

Making it possible to measure frequency, time and distance with an already existing device (the comb part of the optical clock) could potentially reduce the required power consumption and weight, which always come at a premium in space payloads. The microwave and optical frequency transfer has been already demonstrated over fibre with the work presented in this thesis and previous work. Some work has been done on the transfer of microwave frequencies by propagation of an optical frequency comb in free space achieving 1×10^{-15} at 500 s over a 60 m distance Gollapalli and Duan (2010). Despite the short distance, the stability was substantially degraded by the turbulence of the air. In space, where air-turbulence problems are absent, higher levels of stability could potentially be achieved over much longer distances. Much wider frequency combs could also be transferred in space with repect to those transferred over fibre or through the atmosphere as no dispersion effects are found in vacuum. Time transfer has not been demonstrated yet but I believe it could be possible to a very high level of accuracy and I will be investigating this in the near future. Ultra accurate distance measurement has been demonstrated by a number of groups, reaching a precision as high as just a few nm Minoshima and Matsumoto (2000); Ye (2004); Schuhler et al. (2006); Joo and Kim (2006); Swann and Newbury (2006); Salvadé et al. (2008); Joo et al. (2008); Coddington et al. (2009); Cui et al. (2011).

8.3 Future work

8.3.1 Time transfer

The ability to deliver to a user a very stable timing signal in the form of a pulse train could potentially allow the dissemination of time with high accuracy. The crucial difference between transferring time rather than frequency is that the time delay between the transmitter and the receiver must be measured. A simple modulation technique (such as, for example, an intensity variation of a pulse every N pulses as shown in Figure 8.3) could be used as a time marker to measure the round trip delay in a similar fashion as in optical time-domain reflectometry (OTDR) techniques, whilst the optical

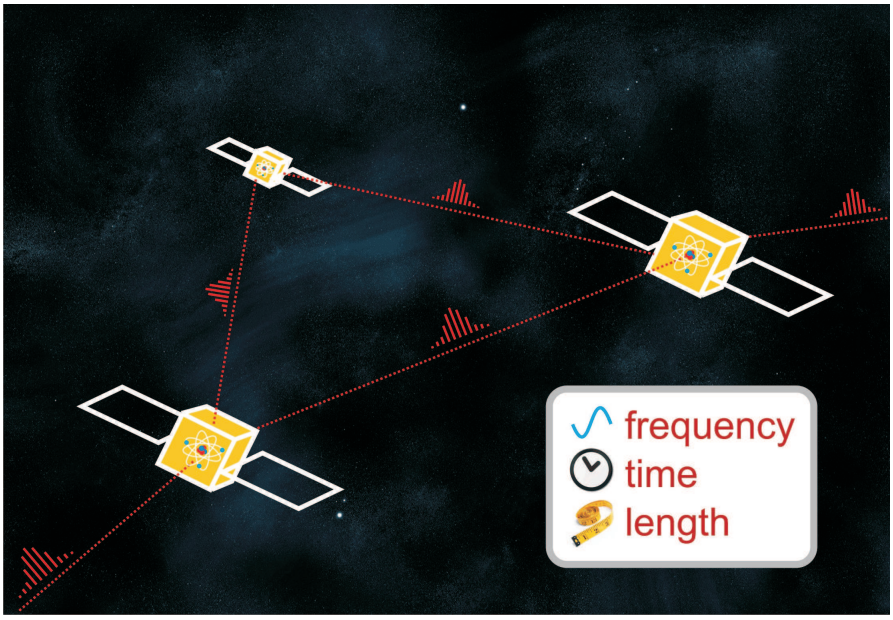

Figure 8.2: If optical clocks are deployed in space, transferring frequency combs rather than single optical carriers could prove advantageous as not only frequency, but also time, can be transferred at the same time and distance measurements between the satellites can also be performed with extremely high accuracy to assist navigation and formation flying. Optical frequency combs are part of an optical clock, so the transfer can be achieved with minimal added complexity, power consumption and weight.

path length changes of the fibre (and hence the time delay) can be stabilized with the same techniques used in frequency transfer experiments. I will be investigating this technique as part of the EMRP-funded project "Accurate time/frequency comparison and dissemination through optical telecommunication networks" already mentioned in Chapter 1.

8.3.2 Spectrum broadening at the remote end

Some potential users might want a stable frequency comb in a different region of the optical spectrum or spanning a wider optical bandwidth than those demonstrated in this thesis. This could be achieved by spectrally broadening the optical spectrum delivered at the user end with a microstructure fibre. In order to achieve spectral broadening, probably a more accurate dispersion compensation than that implemented in the experiments described here would be required. Recompression of the width of pulses to less than 500 fs trasmitted over fibre links up to 120 km has already been shown to be possible by inserting a combination of DCF and dispersion shifting fibre (Yoshida et al., 1998), or prechirping the pulse (Yamamoto and Nakazawa, 2001), or with the help of a programmable pulse shaper in addition to a DCF module (Jiang et al., 2005). With optical amplifiers typically delivering more than 50 mW, it should be possible to induce

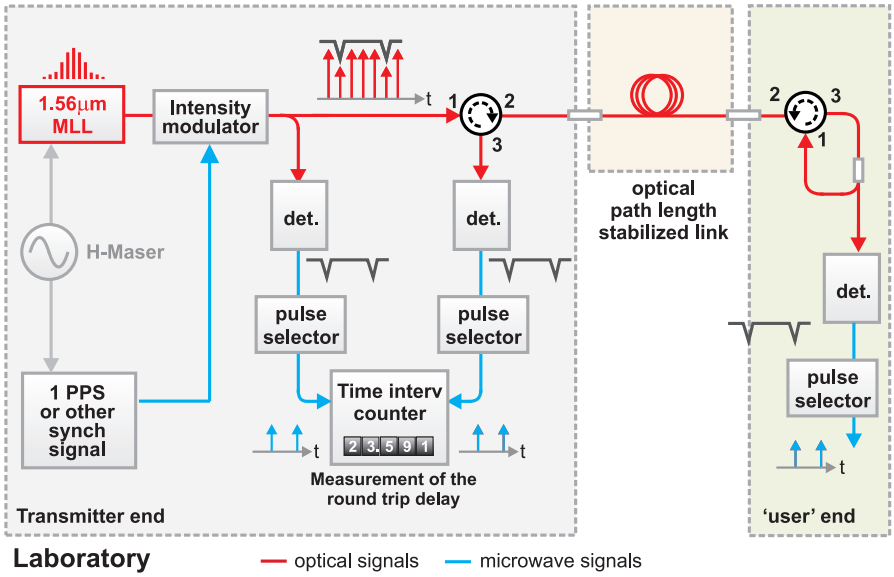

Figure 8.3: Block diagram illustrating a possible solution the time transfer principle by simple intensity modulation of a pulse train. A pulse every N pulses is reduced in power as an it event marker. The event marker can be recovered at the local and user end with a simple photodiode. Any other chosen pulse after the time marker can be used as the *time* marker, benefiting from the high slew rate of the detected pulse. After one-off delay calibration of any uncommon paths at the local and user end, the time delay between between them can be measured with high accuracy and it is maintained constant by the fibre noise cancellation loop.

the self-phase modulation process required for spectral broadening if sub-ps pulses are used.

8.3.3 Final remarks

Often frequency metrology experiments take some time to propagate to other areas of science and industry. This is often because the outstanding performance delivered by the ultra precise frequency metrology experiments is orders of magnitude better than most applications require at each point in time. Frequency metrology experiments are not usually undertaken to solve a specific application need, but they stem from the desire to always set new limits in physics. Most often, the role of frequency metrology becomes that of enabling scientists to dream up new experiments that become feasible once new levels of precision are achieved. The tests of fundamental physics or the geophysics experiments discussed in Chapter 1 are a good example of this. Also, the future practical applications are often very difficult to imagine, as it was certainly difficult to imagine the GPS system when the first atomic clocks were developed.

The same rationale applies to the dissemination of ultra-stable optical frequency combs. Making it possible for scientists in diverse science areas to access a comb of extremely

accurate frequencies delivered from metrological institutes could enable them to envisage new solutions to existing problems or devise new experiments. Hopefully the work presented in this thesis will give a small contribution to help this process.

Appendix A

Fibre stretcher

I constructed the fibre stretcher by tightly winding 250 turns of SMF-28 on a 5-cm diameter piezo tube. Unfortunately, no specifications were available for the piezo tube. The sample is approximately ten years old and no record of its physical characteristics was found. Therefore no accurate calculation of the maximum stretch imposed to the fibre could be performed. However, the optical path length compensation range of the fibre stretcher could be measured directly with the set-up shown in figure A.1.

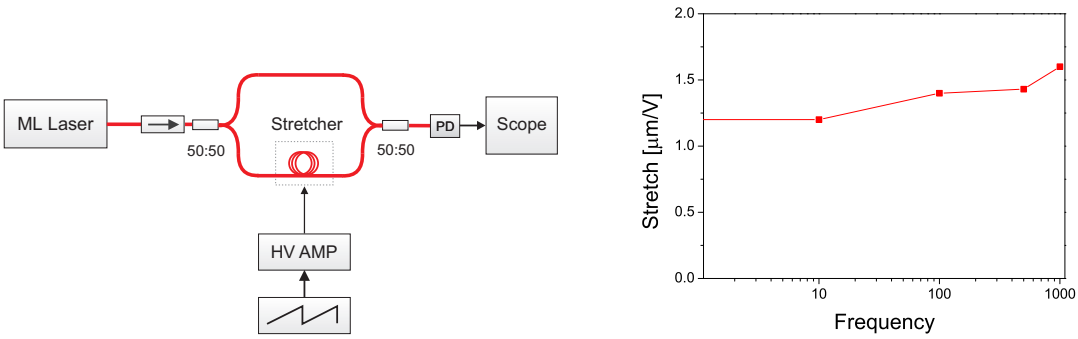

Figure A.1: Stretch as a function of the piezo applied voltage set-up and measurement data. HV AMP: high voltage amplifier

A low frequency triangular voltage is applied at the input of the high voltage amplifier and its amplitude is set to achieve at least 100 interferometric cycles. The output voltage of the photodetector is measured by a digital oscilloscope and the number of optical 2π cycles produced by the change in the voltage applied to the piezo is counted. The same measurement is repeated for a number of frequencies of the driving voltage in order to characterize the frequency response of the fibre stretcher. For this measurement, a fast high voltage amplifier with a maximum output voltage of 150V was used. For an applied voltage range of 130 V, the average number of 2π cycles was 150 over the frequency range 1 Hz to 10 Hz. The stretch will then be $150 \times (1.56/1.48)\,\mu m = 158\,\mu m$, where 1.48 is the refractive index of the core of the SMF-28. The stretch change for a given change

of the applied voltage is

$$\frac{\Delta L}{\Delta V} = 1.2\,\mu\text{m/V} \tag{A.0.1}$$

The maximum applicable voltage to the piezo is 1000 V. Assuming a linear dependence of the stretch to the voltage applied, the maximum stretch achievable is

$$\Delta L = 1.2\,\text{mm} \tag{A.0.2}$$

The thermal coefficient of delay of a bare single mode fibre is approximately 7 $\times 10^{-6}$ K^{-1} (Tateda et al., 1980; K.-C. Lin and Lee, 2004).

The requirements for a fast high voltage amplifier capable of driving the high capacitance of the piezo used in this experiment (10 nF) are relatively demanding. In order to drive the piezo with, say, 1 kHz bandwidth and an output voltage of 1000 V, the amplifier will have to be capable of producing a peak current of

$$I_{\text{peak}} = 2\pi f C V_{\text{pp}} = 63\,\text{mA} \tag{A.0.3}$$

where f is the maximum driving frequency desired, C is the capacitance of the piezo and V_{pp} is the peak-to-peak voltage applied to the piezo.

Appendix B

Single sideband modulator (SSB)

Figure B.1 shows the single sideband modulator used in the fibre transfer experiments described in chapter 5 and 6. The suppression of one of the two sidebands is achieved by preserving a specific phase relationship between the combined outputs of the mixers. The mathematical description of the sideband suppression mechanism goes beyond the scope of this appendix and can be found in most books on modulation techniques. The sine and cosine signals applied at the IF port of each mixer were obtained using a quadrature digital direct synthesiser (DDS). More than 25 dB of sideband suppression was achieved by fine adjustment of both the phase shifter and the amplitude of the signals from the DDS, as shown in figure B.3. An amplifier was used before the power combiner in order to recover some of the power loss due to the mixers and to increase the isolation. Back reflections due to the directional coupler have in fact been shown to compromise significantly the level of sideband suppression achievable. Careful design of the layout and assembly of the single sideband modulator is crucial as any phase instability of this device will be directly superimposed on the useful signal to be measured. The SSB employes a combination of microwave amplifiers, phase shifters and mixers. All these electronics devices exhibit a significant phase sensitivity to temperature, and so it is important that their temperature fluctuations are kept to a minimimum by increasing their thermal mass.

Another possible source of instability is the residual power in the fundamental. Although the SSB suppresses one of the two sidebands, there is still non negligible power at the unshifted frequency (1.5 GHz). When the same harmonics of the repetition rate from two measurement arms are compared using the mixer, a DC voltage is produced due to the residual power into the fundamental. Although the DC voltage can be filtered out with a high pass filter after the mixer, it is preferable to avoid its generation in the first place. Any variation of the DC voltage could change the operating point of the mixer's diodes and therefore also change the power and the phase of the beat signal to be measured. The reduction of power in the fundamental is achieved by implementing a carrier suppression scheme. A fraction of the input signal to the SSB is split off with

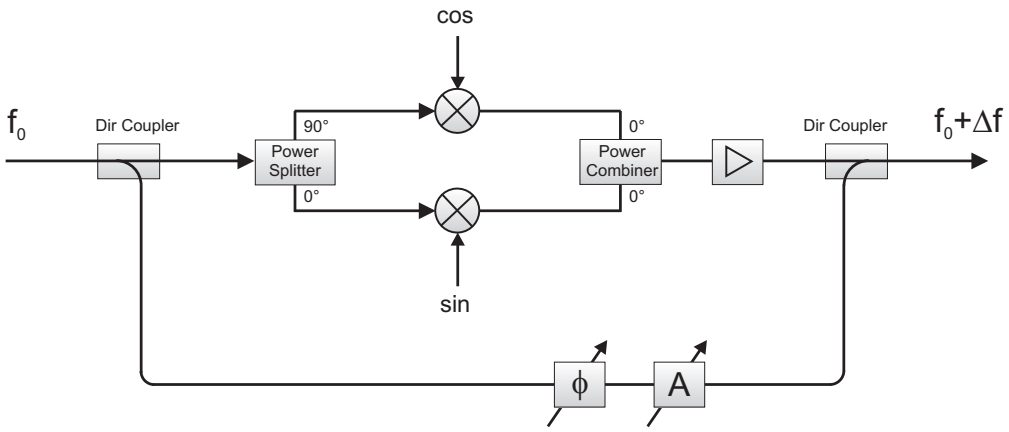

Figure B.1: Set-up of the single sideband modulator. By adjusting the phase shifter and the variable attenuator of the bottom arm, the residual carrier can be suppressed by more than 45 dB. The amplifier increases the isolation making the the sideband and carrier suppression less sensitive to back reflections from the directional coupler.

a directional coupler and then reinjected at the output. If the amplitude of the injected signal is the same as that of the fundamental and the two signals are opposite of phase, then they add destructively. By adjusting the phase shift and the amplitude of the injected signal it was possible to routinely achieve more than 45 dB suppression, as shown in figure B.3. A suppression of 70 dB was possible with careful adjustments.

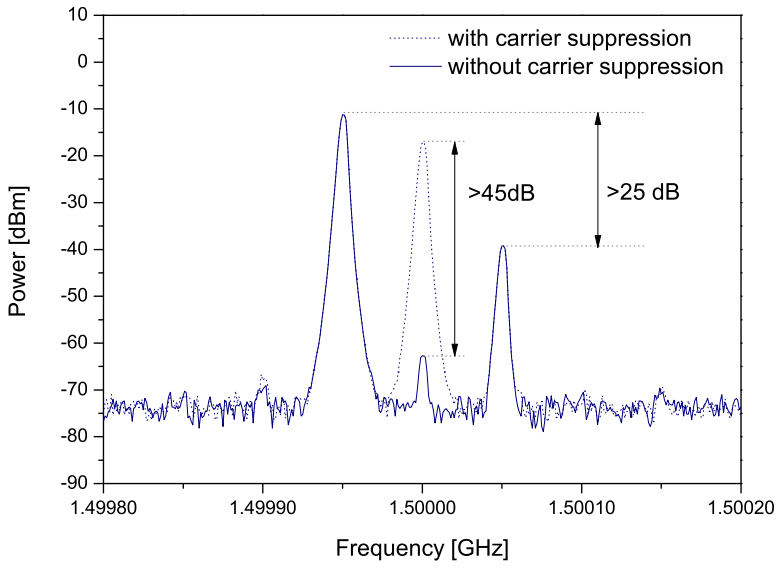

Figure B.2: Interferometric suppression of the unshifted 15^{th} harmonics of the repetition rate

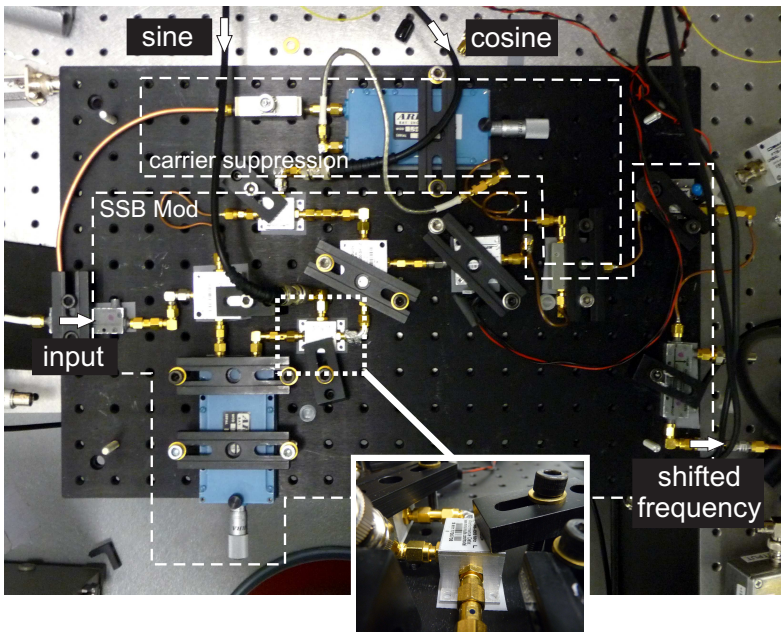

Figure B.3: Photograph of the assembly of the SSB modulator. Inset, the thermal sheet inserted between the microwave components and the breadboard to ensure a good thermal conctact, is visible.

References

Adam, T., Agafonova, N., Aleksandrov, A., Altinok, O., Sanchez, P. A., Anokhina, A., Aoki, S., Ariga, A., Ariga, T., Autiero, D., Badertscher, A., Dhahbi, A. B., Bertolin, A., Bozza, C., Brugiere, T., Brugnera, R., Brunet, F., Brunetti, G., Buontempo, S., Carlus, B., Cavanna, F., Cazes, A., Chaussard, L., Chernyavsky, M., Chiarella, V., Chukanov, A., Colosimo, G., Crespi, M., D'Ambrosio, N., Lellis, G. D., Serio, M. D., Declais, Y., del Amo Sanchez, P., Capua, F. D., Crescenzo, A. D., Ferdinando, D. D., Marco, N. D., Dmitrievsky, S., Dracos, M., Duchesneau, D., Dusini, S., Ebert, J., Efthymiopoulos, I., Egorov, O., Ereditato, A., Esposito, L. S., Favier, J., Ferber, T., Fini, R. A., Fukuda, T., Garfagnini, A., Giacomelli, G., Giorgini, M., Giovannozzi, M., Girerd, C., Goldberg, J., Göllnitz, C., Golubkov, D., Goncharov, L., Gornushkin, Y., Grella, G., Grianti, F., Gschwendtner, E., Guerin, C., Guler, A. M., Gustavino, C., Hagner, C., Hamada, K., Hara, T., Hierholzer, M., Hollnagel, A., Ieva, M., Ishida, H., Ishiguro, K., Jakovcic, K., Jollet, C., Jones, M., Juget, F., Kamiscioglu, M., Kawada, J., Kim, S. H., Kimura, M., Kiritsis, E., Kitagawa, N., Klicek, B., Knuesel, J., Kodama, K., Komatsu, M., Kose, U., Kreslo, I., Lazzaro, C., Lenkeit, J., Ljubicic, A., Longhin, A., Malgin, A., Mandrioli, G., Marteau, J., Matsuo, T., Mauri, N., Mazzoni, A., Medinaceli, E., Meisel, F., Meregaglia, A., Migliozzi, P., Mikado, S., Missiaen, D., Morishima, K., Moser, U., Muciaccia, M. T., Naganawa, N., Naka, T., Nakamura, M., Nakano, T., Nakatsuka, Y., Nikitina, V., Nitti, F., Ogawa, S., Okateva, N., Olchevsky, A., Palamara, O., Paoloni, A., Park, B. D., Park, I. G., Pastore, A., Patrizii, L., Pennacchio, E., Pessard, H., Pistillo, C., Polukhina, N., Pozzato, M., Pretzl, K., Pupilli, F., Rescigno, R., Riguzzi, F., Roganova, T., Rokujo, H., Rosa, G., Rostovtseva, I., Rubbia, A., Russo, A., Sato, O., Sato, Y., Schuler, J., Lavina, L. S., Serrano, J., Sheshukov, A., Shibuya, H., Shoziyoev, G., Simone, S., Sioli, M., Sirignano, C., Sirri, G., Song, J. S., Spinetti, M., Stanco, L., Starkov, N., Stellacci, S., Stipcevic, M., Strauss, T., Takahashi, S., Tenti, M., Terranova, F., Tezuka, I., Tioukov, V., Tolun, P., Tran, N. T., Tufanli, S., Vilain, P., Vladimirov, M., Votano, L., Vuilleumier, J. L., Wilquet, G., Wonsak, B., Wurtz, J., Yoon, C. S., Yoshida, J., Zaitsev, Y., Zemskova, S., and Zghiche, A. (2011). Measurement of the neutrino velocity with the OPERA detector in the CNGS beam.

Antonello, M., Aprili, P., Baibussinov, B., Ceolin, M. B., Benetti, P., Calligarich, E., Canci, N., Carbonara, F., Centro, S., Cesana, A., Cieslik, K., Cline, D. B., Cocco, A.,

Dabrowska, A., Dequal, D., Dermenev, A., Dolfini, R., Farnese, C., Fava, A., Ferrari, A., Fiorillo, G., Gibin, D., Berzolari, A. G., Gninenko, S., Guglielmi, A., Haranczyk, M., Holeczek, J., Ivashkin, A., Kisiel, J., Kochanek, I., Lagoda, J., Mania, S., Mannocchi, G., Menegolli, A., Meng, G., Montanari, C., Otwinowski, S., Periale, L., Piazzoli, A., Picchi, P., Pietropaolo, F., Plonski, P., Rappoldi, A., Raselli, G., Rossella, M., Rubbia, C., Sala, P., Scantamburlo, E., Scaramelli, A., Segreto, E., Sergiampietri, F., Stefan, D., Stepaniak, J., Sulej, R., Szarska, M., Terrani, M., Varanini, F., Ventura, S., Vignoli, C., Wang, H., Yang, X., Zalewska, A., Zaremba, K., Sanchez, P. A., and Serrano, J. (2012). Measurement of the neutrino velocity with the ICARUS detector at the CNGS beam.

Barnes, J. A., Allan, and W., D. (1990). Variances based on data with dead time between the measurements. *U.S. Dept. of Commerce, National Institute of Standards and Technology, Boulder, Colorado*, page 1 v.

Bauch, A., Achkar, J., Bize, S., Calonico, D., Dach, R., R. Hlavac, L. L., Parker, T., Petit, G., Piester, D., Szymaniec, K., and Uhrich, P. (2006). Comparison between frequency standards in Europe and the USA at the 10^{-15} uncertainty level. *Metrologia*, 43(1):109–120.

Bize, S., Diddams, S. A., Tanaka, U., Tanner, C. E., Oskay, W. H., Drullinger, R. E., Parker, T. E., Heavner, T. P., Jefferts, S. R., Hollberg, L., Itano, W., Wineland, D., and Bergquist, J. (2002). Testing the stability of fundamental constants with the $^{199}Hg^{+}$ single-ion optical clock. *Physical Review Letters*, 90(15):4.

Brendel, R., Marianneau, G., and Ubersfeld, J. (1977). Phase and amplitude modulation effects in a phase detector using an incorrectly balanced mixer. *IEEE Trans. Instrum. Meas.*, 26:98–102.

Calhoun, M., Huang, S., and Tjoelker, R. (2007). Stable photonic links for frequency and time transfer in the deep-space network and antenna arrays. *Proceedings of the IEEE*, 95(10):1931 –1946.

Calhoun, M., Sydnor, R., and Diener, W. (2002). A stabilized 100-megahertz and 1-gigahertz reference frequency distribution for Cassini radio science. *Interplanetary Network Progress Rep. 42-148, Jet Propulsion Laboratory, Pasadena, California*, pages 77–85.

CERN (2012). *Press release PR19.11E published on press.web.cern.ch.*

Chen, Y.-F., Jiang, J., and Jones, D. J. (2006). Remote distribution of a mode-locked pulse train with sub 40-as jitter. *Opt. Express*, 14(25):12134–12144.

Chou, C. W., Hume, D. B., Koelemeij, J. C. J., Wineland, D. J., and Rosenband, T. (2010). Frequency comparison of two high-accuracy Al^{+} optical clocks. *Phys. Rev. Lett.*, 104:070802.

Cibiel, G., RBgis, M., Tournier, E., and Llopis, O. (2002). Am noise impact on low level phase noise measurements. *IEEE Trans. Ultrasonics, Ferroelectrics and Freq. Control*, 49(6):784.

Coddington, I., Swann, W. C., Nenadovic, L., and Newbury, N. R. (2009). Rapid and precise absolute distance measurements at long range. *Nature Photonics*, 3(June):351–356.

Cui, M., Zeitouny, M. G., Bhattacharya, N., van den Berg, S. A., and Urbach, H. P. (2011). Long distance measurement with femtosecond pulses using a dispersive interferometer. *Opt. Express*, 19(7):6549–6562.

D. B. Sullivan, D. W. Allan, D. A. H. and Walls, F. L. (1990). Characterization of clocks and oscillators. *NIST Tech Note 1337*.

Datta, S., Joshi, A., and Becker, D. (2009). Low phase noise high power handling InGaAs photodiodes for precise timing applications. *Enabling Photonics Technologies for Defense, Security, and Aerospace Applications, Proceedings of the SPIE,*.

Daussy, C., López, O., Amy-Klein, A., Goncharov, A. N., Guinet, M., Chardonnet, C., Lours, M., Chambon, D., Bize, S., Clairon, A., Santarelli, G., Tobar, M. E., and Luiten, A. N. (2005). Long-distance frequency dissemination with a resolution of 10^{-17}". *Phys. Rev. Lett.*, 94:203904.

Dawkins, S., Mcferran, J., and Luiten, A. (2007). Considerations on the measurement of the stability of oscillators with frequency counters. *Ultrasonics, Ferroelectrics and Frequency Control, IEEE Transactions on*, 54(5):918 –925.

Diddams, S. A., Hollberg, L., Ma, L.-S., and Robertsson, L. (2002). Femtosecond-laser-based optical clockwork with instability$\leq$ less than or equal to 6.3×10^{-16} in 1 s. *Opt. Lett.*, 27(1):58–60.

Diddams, S. A., Hollberg, L., and Mbele, V. (2007). Molecular fingerprinting with the resolved modes of a femtosecond laser frequency comb. *Nature*, 445(7128):627–630.

Dubè, P., Madej, A., Bernard, J., Marmet, L., and Shiner, A. (2009). A narrow linewidth and frequency-stable probe laser source for the $^{88}Sr^{+}$ single ion optical frequency standard. *Applied Physics B: Lasers and Optics*, 95:43–54. 10.1007/s00340-009-3390-6.

Dzuba, V. A. and Flambaum, V. V. (2000). Atomic optical clocks and search for variation of the fine-structure constant. *Phys. Rev. A*, 61:034502.

Dzuba, V. A., Flambaum, V. V., and Webb, J. K. (1999). Calculations of the relativistic effects in many-electron atoms and space-time variation of fundamental constants. *Phys. Rev. A*, 59:230–237.

Ell, R., Morgner, U., Kärtner, F. X., Fujimoto, J. G., Ippen, E. P., Scheuer, V., Angelow, G., Tschudi, T., Lederer, M. J., Boiko, A., and Luther-Davies, B. (2001). Generation of 5-fs pulses and octave-spanning spectra directly from a Ti:sapphire laser. *Opt. Lett.*, 26(6):373–375.

Essen, L. and Parry, J. V. L. (1955). An atomic standard of frequency and time interval: A caesium resonator. *Nature*, 176:280–282.

Fischer, M., Kolachevsky, N., Zimmermann, M., Holzwarth, R., Udem, T., Hänsch, T. W., Abgrall, M., Grünert, J., Maksimovic, I., Bize, S., Marion, H., Santos, F. P. D., Lemonde, P., Santarelli, G., Laurent, P., Clairon, A., Salomon, C., Haas, M., Jentschura, U. D., and Keitel, C. H. (2004). New limits on the drift of fundamental constants from laboratory measurements. *Phys. Rev. Lett.*, 92:230802.

Fortier, T. M., Jones, D. J., and Cundiff, S. T. (2003). Phase stabilization of an octave-spanning Ti:sapphire laser. *Opt. Lett.*, 28(22):2198–2200.

Fujieda, M., Kumagai, M., Nagano, S., Yamaguchi, A., Hachisu, H., and Ido, T. (2011). All-optical link for direct comparison of distant optical clocks. *Opt. Express*, 19(17):16498–16507.

Gill, P. (2011). When should we change the definition of the second? *Philosophical transactions. Series A, Mathematical, physical, and engineering sciences*, 369(1953):4109–4130.

Gjestland, T. T. (2008). Background noise levels in europe. *SINTEF A6631*.

Gokhale (2006). Deployment of fiber optic networks through underground sewers in north america. *Journal of Transportation Engineering*, 132(8):672–682.

Gollapalli, R. and Duan, L. (2010). Atmospheric timing transfer using a femtosecond frequency comb. *Photonics Journal, IEEE*, 2(6):904 –910.

Grosche, G., Terra, O., Predehl, K., Holzwarth, R., Lipphardt, B., Vogt, F., Sterr, U., and Schnatz, H. (2009). Optical frequency transfer via 146 km fiber link with 10^{-19} relative accuracy. *Opt. Lett.*, 34(15):2270–2272.

Guéna, J., Abgrall, M., Rovera, D., Laurent, P., Chupin, B., Lours, M., Santarelli, G., Rosenbusch, P., Tobar, M. E., Li, R., Gibble, K., Clairon, A., , and Bize, S. (2012). Progress in atomic fountains at lne-syrte. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, (59):391–410.

Hartnett, J. G., Nand, N. R., and Lu, C. (2012). Ultra-low-phase-noise cryocooled microwave dielectric-sapphire-resonator oscillators. *Applied Physics Letters*, 100(18).

He, Y., Hsu, M., Gray, M., Wouters, M., Warrington, R., Luiten, A., Baldwin, K., Shaddock, D., Orr, B., and Aben, G. (2011). An optical fibre-based frequency dissemination network for Australia. *2011 Joint Conference of the IEEE International*

Frequency Control and the European Frequency Symposium and Time Forum (IFC-S/EFTF), pages 1–3.

Holman, K. W., Hudson, D. D., Ye, J., and Jones, D. J. (2005). Remote transfer of a high-stability and ultralow-jitter timing signal. *Opt. Lett.*, 30(10):1225–1227.

Holman, K. W., Jones, D. J., Hudson, D. D., and Ye, J. (2004). Precise frequency transfer through a fiber network by use of 1.5-μm mode-locked sources. *Opt. Lett.*, 29(13):1554–1556.

Holman, K. W., Jones, R. J., Marian, A., Cundiff, S. T., and Ye, J. (2003). Intensity-related dynamics of femtosecond frequency combs. *Opt. Lett.*, 28(10):851–853.

Holzwarth, R., Udem, T., Hänsch, T. W., Knight, J. C., Wadsworth, W. J., and Russell, P. S. J. (2000). Optical frequency synthesizer for precision spectroscopy. *Phys. Rev. Lett.*, 85:2264–2267.

Hou, D., Li, P., Liu, C., Zhao, J., and Zhang, Z. (2011). Long-term stable frequency transfer over an urban fiber link using microwave phase stabilization. *Opt. Express*, 19(2):506–511.

Ivanov, E. N., Diddams, S. A., , and Hollberg, L. (2005). Study of the excess noise associated with demodulation of ultra-short infrared pulses. *IEEE Trans. Ultrasonics, Ferroelectrics and Freq. Control*, 52:1068–1074.

Ivanov, E. N., Hollberg, L., and Diddams, S. A. (2003). Analysis of noise mechanisms limiting the frequency stability of microwave signals generated with a femtosecond laser. *IEEE J. Sel. Top. Quantum Electon.*, 9:1059–1065.

Ivanov, E. N., McFerran, J. J., Diddams, S. A., and Hollberg, L. (2007). Noise properties of microwave signals synthesized with femtosecond lasers. *IEEE Trans. Ultrason. Ferroelectr. Freq. Contr.*, 54:736–745.

Jefferts, S., Heavner, T., Parker, T., and J.H.Shirley (2007). NIST cesium fountains current status and future prospects. *Time and Frequency Metrology. Edited by Jones, R. Jason. Proceedings of the SPIE*, 6673:6673–08.

Jiang, H., Kéfélian, F., Crane, S., Lopez, O., Lours, M., Millo, J., Holleville, D., Lemonde, P., Chardonnet, C., Amy-Klein, A., and Santarelli, G. (2008). Long-distance frequency transfer over an urban fiber link using optical phase stabilization. *J. Opt. Soc. Am. B*, 25(12):2029–2035.

Jiang, Y. Y., Ludlow, A. D., Lemke, N. D., Fox, R. W., Sherman, J. A., Ma, L. S., and Oates, C. W. (2011). Making optical atomic clocks more stable with 10^{-16} -level laser stabilization,. *Nature Photonics*, 5:158–161.

Jiang, Z., Yang, S.-D., Leaird, D. E., and Weiner, A. M. (2005). Fully dispersion-compensated ∼500 fs pulse transmission over 50 km single-mode fiber. *Opt. Lett.*, 30(12):1449–1451.

Jones, D. J., Diddams, S. A., Ranka, J. K., Stentz, A., Windeler, R. S., Hall, J. L., and Cundiff, S. T. (2000). Carrier-envelope phase control of femtosecond mode-locked lasers and direct optical frequency synthesis. *Science*, 288(5466):635–639.

Joo, K.-N. and Kim, S.-W. (2006). Absolute distance measurement by dispersive interferometry using a femtosecond pulse laser. *Opt. Express*, 14(13):5954–5960.

Joo, K.-N., Kim, Y., and Kim, S.-W. (2008). Distance measurements by combined method based on a femtosecond pulse laser. *Opt. Express*, 16(24):19799–19806.

K.-C. Lin, C.-J. L. and Lee, W.-Y. (2004). Effects of gamma radiation on optical fibre sensors. *Optoelectronics, IEE Proceedings*, 151(1):12–15.

Kärtner (2012). Status on pulsed timing distribution systems and implementations at desy, fermi and xfel.

Kaye, G. W. C. and Laby, T. H. (1995). Tables of physical and chemical constants and some mathematical functions. *Longman*.

Kéfélian, F., Lopez, O., Jiang, H., Chardonnet, C., Amy-Klein, A., and Santarelli, G. (2009). High-resolution optical frequency dissemination on a telecommunications network with data traffic. *Opt. Lett.*, 34(10):1573–1575.

Kleppner, D. (2007). Time too good to be true. *Physics Today*, 59(3):10–11.

Kumagai, M., Fujieda, M., Nagano, S., and Hosokawa, M. (2009). Stable radio frequency transfer in 114 km urban optical fiber link. *Opt. Lett.*, 34(19):2949–2951.

Lea, S. N. (2008). Limits to time variation of fundamental constants from comparisons of atomic frequency standards. *The European Physical Journal - Special Topics*, 163:37–53.

Li, R., Gibble, K., and Szymaniec, K. (2011). Improved accuracy evaluation of the NPL-CsF2 primary frequency standard.

Lodewyck, J., Westergaard, P. G., Lecallier, A., Lorini, L., and Lemonde, P. (2010). Frequency stability of optical lattice clocks. *New Journal of Physics*, 12(6):065026.

Löhl, F., Arsov, V., Felber, M., Hacker, K., Jalmuzna, W., Lorbeer, B., Ludwig, F., Matthiesen, K.-H., Schlarb, H., Schmidt, B., Schmüser, P., Schulz, S., Szewinski, J., Winter, A., and Zemella, J. (2010). Electron bunch timing with femtosecond precision in a superconducting free-electron laser. *Phys. Rev. Lett.*, 104:144801.

Lopez, O., Amy-Klein, A., Daussy, C., Chardonnet, C., Narbonneau, F., Lours, M., and Santarelli, G. (2008). 86-km optical link with a resolution of 2×10^{-18} for rf frequency transfer. *The European Physical Journal D - Atomic, Molecular, Optical and Plasma Physics*, 48:35–41.

Lopez, O., Amy-Klein, A., Lours, M., Chardonnet, C., and Santarelli, G. (2010). High-resolution microwave frequency dissemination on an 86-km urban optical link. *Applied Physics B*, 98:723–727.

Ludlow, A. D., Huang, X., Notcutt, M., Zanon-Willette, T., Foreman, S. M., Boyd, M. M., Blatt, S., and Ye, J. (2007). Compact, thermal-noise-limited optical cavity for diode laser stabilization at 1×10^{-15}. *Opt. Lett.*, 32(6):641–643.

Lutes, G. (1980). Experimental optical fiber communications link. *The Telecommunications and Data Acquisition Progress Report, TDA PR 42-59, Jul./Aug. 1980*, pages 77–85.

Lutes, G. (1986). Experimental optical fiber communications link. *The Telecommunications and Data Acquisition Progress Report, TDA PR 42-87, Jul./Sep. 1986*, pages 1–9.

Lyons, H. (1952). Spectral lines as frequency standards. *Annals of the New York Academy of Sciences*, 55:831–871.

Ma, L.-S., Bi, Z., Bartels, A., Kim, K., Robertsson, L., Zucco, M., Windeler, R. S., Wilpers, G., Oates, C., Hollberg, L., and Diddams, S. A. (2007). Frequency uncertainty for optically referenced femtosecond laser frequency combs. *IEEE J. Quant. Electron.*, 43:139–46.

Ma, L.-S., Jungner, P., Ye, J., and Hall, J. L. (1994). Delivering the same optical frequency at two places: accurate cancellation of phase noise introduced by an optical fiber or other time-varying path. *Opt. Lett.*, 19(21):1777–1779.

Ma, L. S., Zhiyi, B., Bartels, A., Robertsson, L., Zucco, M., Windeler, R. S., Wilpers, G., Oates, C., Hollberg, L., and Diddams, S. A. (2004). Optical frequency synthesis and comparison with uncertainty at the 10^{-19} level. *Science*, 303:1843–1845.

Markowitz, W., Hall, R. G., Essen, L., and Parry, J. V. L. (1958). Frequency of cesium in terms of ephemeris time. *Phys. Rev. Lett.*, 1:105–107.

Marra, G., Margolis, H. S., Lea, S. N., and Gill, P. (2010). High-stability microwave frequency transfer by propagation of an optical frequency comb over 50 km of optical fiber. *Opt. Lett.*, 35(7):1025–1027.

Marra, G., Margolis, H. S., and Richardson, D. J. (2012). Dissemination of an optical frequency comb over fiber with 3×10^{-18} fractional accuracy. *Opt. Express*, 20(2):1775–1782.

Marra, G., Slavík, R., Margolis, H. S., Lea, S. N., Petropoulos, P., Richardson, D. J., and Gill, P. (2011). High-resolution microwave frequency transfer over an 86-km-long optical fiber network using a mode-locked laser. *Opt. Lett.*, 36(4):511–513.

Millo, J., Magalhães, D. V., Mandache, C., Le Coq, Y., English, E. M. L., Westergaard, P. G., Lodewyck, J., Bize, S., Lemonde, P., and Santarelli, G. (2009). Ultrastable lasers based on vibration insensitive cavities. *Phys. Rev. A*, 79:053829.

Minoshima, K. and Matsumoto, H. (2000). High-accuracy measurement of 240-m distance in an optical tunnel by use of a compact femtosecond laser. *Appl. Opt.*, 39(30):5512–5517.

Morgner, U., Ell, R., Metzler, G., Schibli, T. R., Kärtner, F. X., Fujimoto, J. G., Haus, H. A., and Ippen, E. P. (2001). Nonlinear optics with phase-controlled pulses in the sub-two-cycle regime. *Phys. Rev. Lett.*, 86:5462–5465.

Musha, M., Hong, F.-L., Nakagawa, K., and Ueda, K. (2008). Coherent optical frequency transfer over 50-km physical distance using a 120-km-long installed telecom fiber network. *Opt. Express*, 16(21):16459–16466.

Notcutt, M., Ma, L.-S., Ludlow, A. D., Foreman, S. M., Ye, J., and Hall, J. L. (2006). Contribution of thermal noise to frequency stability of rigid optical cavity via hertz-linewidth lasers. *Phys. Rev. A*, 73:031804.

Olivier, L., Haboucha, A., Chanteau, B., Chardonnet, C., Amy-Klein, A., and Santarelli, G. (2012). Ultra-stable long distance optical frequency distribution using the internet fiber network.

Parker, B., Marra, G., Lea, S., Margolis, H., Bayvel, P., and Gill, P. (2012). *in preparation.*

Peik, E., Lipphardt, B., Schnatz, H., Schneider, T., Tamm, C., and Karshenboim, S. G. (2004). New limit on the present temporal variation of the fine structure constant. *Physical Review Letters*, 93(17):4.

Predehl, K., Grosche, G., Raupach, S. M. F., Droste, S., Terra, O., Alnis, J., Legero, T., Hänsch, T. W., Udem, T., Holzwarth, R., and Schnatz, H. (2012). A 920-kilometer optical fiber link for frequency metrology at the 19th decimal place. *Science*, 336.

Prestage, J. D., Tjoelker, R. L., and Maleki, L. (1995). Atomic clocks and variations of the fine structure constant. *Phys. Rev. Lett.*, 74:3511–3514.

Primas, L., Logan, R.T., J., and Lutes, G. (1989). Applications of ultra-stable fiber optic distribution systems. *Frequency Control, 1989., Proceedings of the 43rd Annual Symposium on.*

Primas, L., Lutes, G., and Sydnor, R. (1988). Fiber optic frequency transfer link. *Proceedings of the 42nd Annual Frequency Control Symposium, 1988.,*, pages 478–484.

Ramond, T. M., Diddams, S. A., Hollberg, L., and Bartels, A. (2002). Phase-coherent link from optical to microwave frequencies by means of the broadband continuum from a 1-GHz Ti:sapphire femtosecondoscillator. *Opt. Lett.*, 27(20):1842–1844.

Reichert, J., Holzwarth, R., Udem, T., and Hänsch, T. (1999). Measuring the frequency of light with mode-locked lasers. *Optics Communications*, 172(16):59 – 68.

Rosenband, T., Hume, D. B., Schmidt, P. O., Chou, C. W., Brusch, A., Lorini, L., Oskay, W. H., Drullinger, R. E., Fortier, T. M., Stalnaker, J. E., Diddams, S. A., Swann, W. C., Newbury, N. R., Itano, W. M., Wineland, D. J., and Bergquist, J. C. (2008). Frequency ratio of Al^+ and Hg^+ single-ion optical clocks; metrology at the 17^{th} decimal place. *Science*, 319(5871):1808–1812.

Rubiola, E. (2005). On the measurement of frequency and of its sample variance with high-resolution counters. *Review of Scientific Instruments*, 76(5):054703.

Rutman, J. (1978). Characterization of phase and frequency instabilities in precision frequency sources: fifteen years of progress. *Proc. IEEE*, 66:1048–1075.

Salvadé, Y., Schuhler, N., Lévêque, S., and Floch, S. L. (2008). High-accuracy absolute distance measurement using frequency comb referenced multiwavelength source. *Appl. Opt.*, 47(14):2715–2720.

Sandberg, U. (2005). The multi-coincidence peak around 1000 Hz in tyre-road noise spectra. *Proceedings of Euronoise, Naples, Italy.*

Schnatz, H., Lipphardt, B., Helmcke, J., Riehle, F., and Zinner, G. (1996). First phase-coherent frequency measurement of visible radiation. *Phys. Rev. Lett.*, 76:18–21.

Schuhler, N., Salvadé, Y., Lévêque, S., Dändliker, R., and Holzwarth, R. (2006). Frequency-comb-referenced two-wavelength source for absolute distance measurement. *Opt. Lett.*, 31(21):3101–3103.

Stoehr, H., Mensing, F., Helmcke, J., and Sterr, U. (2006). Diode laser with 1 Hz linewidth. *Opt. Lett.*, 31(6):736–738.

Sullivan, D., Allan, D., Howe, D., and Walls, F. (1990). Characterization of clocks and oscillators. *NIST Tech Note 1337.*

Svehla, D. (2008). A novel design for the navigation system and proposal and the proposal to unify the timing and the positioning system using giove follow-on. *presentation at the Workshop on ACES and future GNSS-based Earth observation and navigation, Munich, Germany.*

Swann, W. C. and Newbury, N. R. (2006). Frequency-resolved coherent lidar using a femtosecond fiber laser. *Opt. Lett.*, 31(6):826–828.

Tamura, K., Ippen, E. P., Haus, H. A., and Nelson, L. E. (1993). 77-fs pulse generation from a stretched-pulse mode-locked all-fiber ring laser. *Opt. Lett.*, 18(13):1080–1082.

Tapley, B. D., Bettadpur, S., Watkins, M., and Reigber, C. (2004). The gravity recovery and climate experiment: mission overview and early results. *Geophys. Res. Lett.*, 31(L09607).

Tateda, M., Tanaka, S., and Sugawara, Y. (1980). Thermal characteristics of phase shift in jacketed optical fibers. *Applied Optics*, 19(5):770–773.

Taylor, J., Datta, S., Hati, A., Nelson, C., Quinlan, F., Joshi, A., and Diddams, S. (2011). Characterization of power-to-phase conversion in high-speed p-i-n photodiodes. *Photonics Journal, IEEE*, 3(1):140 –151.

Terra, O., Grosche, G., and Schnatz, H. (2010). Brillouin amplification in phase coherent transfer of optical frequencies over 480 km fiber. *Opt. Express*, 18(15):16102–16111.

Thompson, D. (2008). Railway noise and vibration: mechanisms, modelling and means of control. *Elsevier Science*.

Vig, J. R. (1999). IEEE standard definitions of physical quantities for fundamental frequency and time metrology-random instabilities (IEEE Standard 1139-1999). *IEEE, New York*.

Walton, B., Margolis, H., Tsatourian, V., and Gill, P. (2008). Transportable optical frequency comb based on a mode-locked fibre laser. *Optoelectronics, IET*, 2(5):182 –187.

Watabe, K., Hartnett, J. G., Locke, C. R., Santarelli, G., Yanagimachi, S., Shimazaki, T., Ikegami, T., and ichi Ohshima, S. (2006). Short term frequency stability tests of two cryogenic sapphire oscillators. *Japanese Journal of Applied Physics*, 45(12):9234–9237.

Webster, S. and Gill, P. (2011). Force-insensitive optical cavity. *Opt. Lett.*, 36(18):3572–3574.

Webster, S. A., Oxborrow, M., Pugla, S., Millo, J., and Gill, P. (2008). Thermal-noise-limited optical cavity. *Phys. Rev. A*, 77:033847.

Weyers, S., Gerginov, V., Nemitz, N., Li, R., and Gibble, K. (2012). Distributed cavity phase frequency shifts of the caesium fountain ptb-csf2. *Metrologia*, 49(1):82.

Williams, P. A., Swann, W. C., and Newbury, N. R. (2008). High-stability transfer of an optical frequency over long fiber-optic links. *J. Opt. Soc. Am. B*, 25(8):1284–1293.

Yamamoto, T. and Nakazawa, M. (2001). Third- and fourth-order active dispersion compensation with a phase modulator in a terabit-per-second optical time-division multiplexed transmission. *Opt. Lett.*, 26(9):647–649.

Yao, X. S., Lutes, G., Logan, R. T. J., and Maleki, L. (1994). Field demonstration of X-band photonic antenna remoting in the deep space network. *The Telecommunications and Data Acquisition Progress Report*, pages 29–34.

Ye, J. (2004). Absolute measurement of a long, arbitrary distance to less than anoptical fringe. *Opt. Lett.*, 29(10):1153–1155.

Yoshida, E., Yamamoto, T., Sahara, A., and Nakazawa, M. (1998). 320 Gbit/s TDM transmission over 120 km using 400 fs pulse train. *Electronics Letters*, 34(10):1004 –1005.

Young, B. C., Cruz, F. C., Itano, W. M., and Bergquist, J. C. (1999). Visible lasers with subhertz linewidths. *Phys. Rev. Lett.*, 82:3799–3802.

Zhang, L., Chang, L., Dong, Y., Xie, W., He, H., and Hu, W. (2011). Phase drift cancellation of remote radio frequency transfer using an optoelectronic delay-locked loop. *Opt. Lett.*, 36(6):873–875.

Zhang, W., Xu, Z., Lours, M., Boudot, R., Kersale, Y., Santarelli, G., and Le Coq, Y. (2010). Sub-100 attoseconds stability optics-to-microwave synchronization. *AApplied Physics Letters*, 96(21):211105.

Publications

High-stability microwave frequency transfer by propagation of an optical frequency comb over 50 km of optical fiber

Giuseppe Marra,* Helen S. Margolis, Stephen N. Lea, and Patrick Gill

National Physical Laboratory, Hampton Road, Teddington, Middlesex TW11 0LW, UK
**Corresponding author: giuseppe.marra@npl.co.uk*

Received November 24, 2009; revised February 15, 2010; accepted February 15, 2010; posted February 24, 2010 (Doc. ID 120478); published March 30, 2010

A high-stability microwave frequency reference is transferred over 50 km of spooled optical fiber by propagation of a 90-nm-wide optical frequency comb centered at 1.56 μm. Environmentally induced fiber phase noise is actively suppressed by returning the optical frequency comb via a separate fiber. The stability of the microwave frequency delivered is measured at the "user" end, eliminating the need for the assumption of correlation between the noise in the forward and return paths. The measured transfer stability is $4.6 \times 10^{-15}\tau^{-1}$, rendering this technique suitable for the transfer of state-of-the-art frequency references.

OCIS codes: 120.3940, 140.4050, 060.2360, 060.5530.

The transfer of ultrastable microwave and optical frequencies over optical fiber links has recently become an area of active research, driven primarily by the need to compare state-of-the-art atomic frequency standards in distantly located laboratories. Currently, remote frequency comparison is usually achieved by satellite-based techniques such as two-way satellite time and frequency transfer or GPS carrier phase observations [1], but these are unable to preserve the stability delivered by state-of-the-art microwave and optical frequency standards [2]. In contrast, recent experiments on optical fiber-based techniques have demonstrated the transfer of microwave and optical frequencies with a stability a few orders of magnitude higher than satellite-based techniques. Most of this work has focused on either the transfer of a microwave frequency by modulation of an optical carrier [3,4] or the transfer of an optical frequency by direct transmission of the optical carrier itself [3,5,6]. The best transfer stability achieved using these techniques is, respectively, 1.3×10^{-15} at 1 s over a distance of 86 km [4] and better than a few parts in 10^{-16} at 1 s over fiber lengths spanning from 86 to 251 km [5–8].

A third technique, whereby optical and microwave references are transferred at the same time by propagation of an optical frequency comb, has been investigated by only one other group. A microwave transfer stability of better than $9 \times 10^{-15}\tau^{-1/2}$ was demonstrated over a 6.9 km round-trip installed fiber link, using separate fibers for the forward and return paths [9]. However, the transfer stability was measured at the transmitter end, after the comb had traveled a round trip. The validity of this approach relies on the assumption of sufficient correlation between the fiber-induced phase changes in the forward and return paths. Insufficient correlation can degrade the actual stability delivered at the user end, which cannot be detected after a round trip. In this work, we transfer an optical frequency comb over a considerably longer round-trip distance of 100 km and measure the true microwave transfer stability at the user end. An improved transfer stability is reported, scaling approximately as τ^{-1} rather than $\tau^{-1/2}$. In addition, the transferred 3 dB optical bandwidth is ~90 nm instead of 5–15 nm, corresponding to over 10^5 optical modes being transmitted.

The experimental setup is shown in Fig. 1. A commercial 1.56 μm amplified erbium-doped mode-locked fiber laser emits sub-150-fs optical pulses at a repetition rate of 100 MHz. The average optical power is approximately 65 mW, required to detect the repetition rate with a sufficient signal-to-noise ratio (SNR) after 100 km of the fiber. The repetition rate is phase-locked to a hydrogen maser-referenced radio frequency (RF) synthesizer with a locking bandwidth of a few hundred hertz. The output of the mode-locked laser is split into two optical paths by a 90:10 splitter. After passing through an optical isolator, the high power beam is transmitted to the user end via

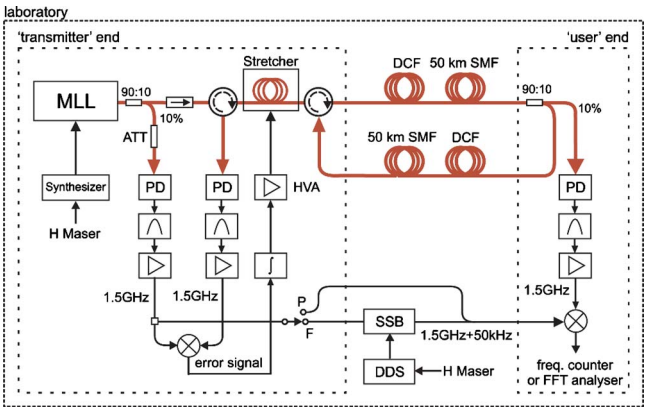

Fig. 1. (Color online) Block diagram of the experimental setup. The noise introduced by the fiber is measured by comparing the local and returned repetition rates and suppressed with a fiber stretcher. A simple connection change switches between phase noise (P) and frequency stability (F) measurements of the repetition rate signal at the user end: MLL, mode-locked laser; PD, photodetector; DCF, dispersion-compensating fiber; SSB, single-sideband modulator; HVA, high-voltage amplifier; H Maser, hydrogen maser; DDS, direct digital synthesizer.

50 km of a standard single-mode fiber (SMF-28). Here, another 90:10 splitter returns the larger fraction of the optical power back to the "transmitter" end via a second 50 km fiber. This arrangement has been chosen to replicate a typical fiber network, where installed unidirectional optical amplifiers prevent the use of the same fiber for the forward and return optical paths. Chromatic dispersion broadens the laser pulses to ~150 ns over 100 km of SMF-28. By inserting a dispersion-compensating fiber (DCF) module in each 50 km arm (dispersion of −864 ps/nm at 1560 nm and loss of 3.8 dB), the pulse duration is recompressed to less than 220 ps. Although this is still 3 orders of magnitude higher than the initial pulse duration, the large ratio between the repetition period and the pulse width was sufficient to achieve a SNR of the extracted repetition rate, after both 50 and 100 km, in excess of 80 dB (resolution bandwidth of 1 kHz).

The low power output of the 90:10 power splitter at the transmitter end provides the low-noise reference against which the round-trip phase noise introduced by the 100 km link is measured. The laser repetition rate is detected here (local repetition rate), after a round trip (returned repetition rate), and at the user end (remote repetition rate) using battery-powered high-speed GaAs photodetectors. In each case the 15th harmonic of the repetition rate (1.5 GHz) is selected with a narrow tunable bandpass filter that suppresses the nearest harmonics by more than 30 dB. The RF power at the filter outputs is approximately −21, −40, and −47 dBm for the local, remote, and returned repetition rate signals, respectively. Two low-noise cascaded microwave amplifiers are employed for the local repetition rate and three for the remote and returned repetition rates to reach efficient driving levels of the subsequent microwave devices.

The error signal constructed by comparing the local and returned repetition rates with a microwave mixer is amplified, integrated, and fed to a high voltage amplifier that controls a fiber stretcher. This compensates for the measured phase noise within the feedback bandwidth, which arises mainly from mechanical and thermal perturbations of the optical path length in the fiber, with secondary contributions from other effects such as optical nonlinearities and fluctuations in chromatic or polarization-mode dispersion. The feedback bandwidth is limited by the high voltage amplifier to ~10 Hz. With improved hardware it could be increased nearer to the theoretical limit of 2 kHz, imposed by the 0.5 ms round-trip travel time of the pulse train. The fiber stretcher was built in-house by winding 250 turns of standard SMF-28 fiber onto a 5 cm diameter piezo tube and can be used to compensate for changes of up to 1.7 mm in the optical path length. Two optical circulators, one preceding and one following the fiber stretcher, were used to separate the pulses traveling in the forward and backward directions. Achieving low levels of reflected power at the input of the photodetector extracting the returned repetition rate is crucial if no optical amplifiers are used along the optical link. The useful optical power available after a round trip is reduced by over 30 dB due to the cumulative losses of the 100 km of SMF-28, DCF, and optical circulators, becoming of a comparable level to the unwanted reflections caused by optical interconnections. Acceptable levels of reflections are achieved in our setup by fusion-splicing each end of the fiber stretcher to the optical circulators and by using angle-polished connectors everywhere else.

The mechanical and thermal stabilities of the microwave electronics are also important. In particular, insufficient thermal stability of the microwave mixer used to construct the error signal for phase noise cancellation can result in increased transfer instability beyond a few seconds, with any unwanted variations of the mixer output voltage being converted directly into a phase error on the extracted repetition rate. All the microwave components are therefore rigidly mounted onto aluminum breadboards. The optical power splitters, isolator, and circulators, which also show non-negligible sensitivity to changes in ambient temperature, are enclosed in a box lined with thermal foam sheets.

The phase noise and frequency stability of the repetition rate delivered to the user are measured with a minimal change to the setup. For the phase noise measurement, the remote repetition rate is phase compared with the local repetition rate using a microwave mixer and processed with a fast Fourier transform analyzer. To measure the fractional frequency stability, a small frequency offset is added to the 15th harmonic of the local repetition rate by inserting an SSB modulator driven by a digital direct synthesizer, whose clock is referenced to a hydrogen maser. The residual noise floor of the SSB is measured to be $8\times10^{-16}\tau^{-1}$. An offset frequency of 50 kHz is chosen as this falls into the range where the Λ-type Agilent 53132A frequency counter [10] used in this work shows the best resolution. The power levels applied at the input of the first amplifier and the amplifiers' noise figure impose white phase noise floors of −137 and −130 dBc/Hz on the remote and returned repetition rates, respectively. However, the self-homodyne detection of the laser repetition rate noise due to the delay introduced by the fiber exceeds the amplifier noise for frequencies above a few tens of hertz. This noise would be substantially lower if the transfer technique described were used to transfer a state-of-the-art microwave frequency reference such as the best commercial room temperature oscillators or research-type cryogenic sapphire oscillators, or if the comb were referenced to an ultra-stable optical frequency reference such as those based on high-finesse cavities [11,12].

The phase noise measured with and without active phase noise suppression is shown in Fig. 2. At 1 Hz offset from the carrier the noise is suppressed by more than 10 to −101 dBc/Hz. The transfer frequency stability, as measured at the user end, is shown in Fig. 3. Each data point is obtained from independent frequency stability measurements at different counter gate times, with the measurement time limited to ~250 s by the range of the fiber

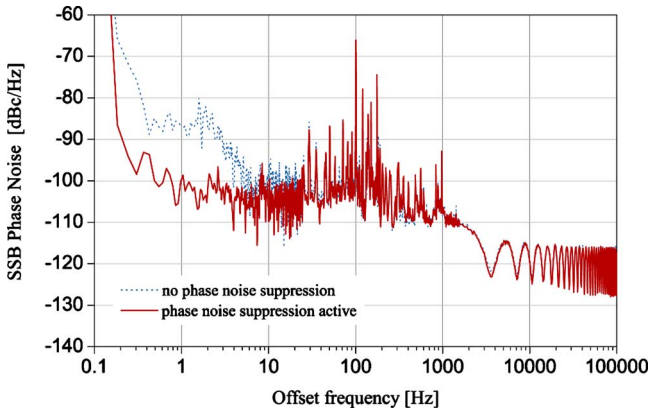

Fig. 2. (Color online) Phase noise of the 15th harmonic of the repetition rate delivered at the user end of the fiber with and without phase-noise suppression. The dips above a few kilohertz are due to self-homodyne detection of the laser noise. Their frequencies correspond to multiples of the inverse of the pulse propagation time over the fiber length.

stretcher. Full compensation of thermally induced optical path length changes could be achieved in future work by means of a thermally controlled fiber spool inserted in series with the fiber stretcher. A weighted linear fit to the data, with the slope constrained to be -1, gives a measured transfer instability of $4.6 \times 10^{-15}\tau^{-1}$, demonstrating the suitability of this technique for transferring the stability of most high quality microwave or optical references. The measurement noise floor due to the optical to microwave conversion is measured by replacing the 50 km fiber spools with 2 m fibers and inserting optical attenuators to keep the detected signal levels the same. The residual instability due to the microwave electronics following the photodetectors is measured by mimicking the signals from the photodetectors with low-noise frequency synthesizers. The lower instability measured in this case, limited by the SSB, indicates that the amplitude-to-phase noise conversion in the

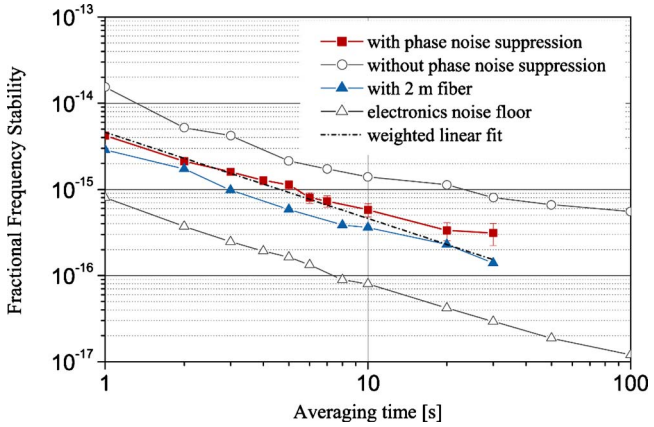

Fig. 3. (Color online) Fractional frequency stability of the 15th harmonic of the repetition rate delivered at the user end of the fiber with (squares) and without (circles) phase noise suppression. The optical noise floor (filled triangles) and the noise floor imposed by the electronics (open triangles) are also shown.

photodetectors [13] or the residual amplitude sensitivity of the microwave mixer, or a combination of both effects, might be responsible for the higher optical noise floor. Amplitude noise measurements performed with an RF power detector on the input signal applied to the mixer do indeed show levels 15 dB higher when the 1.5 GHz signal is obtained from the photodetectors rather than the RF synthesizers.

In summary, the frequency stability of a microwave reference transferred by propagation of an optical frequency comb over 50 km of a laboratory-based fiber network has been measured directly at the user end and is found to be $4.6 \times 10^{-15}\tau^{-1}$. In the future we plan to test this frequency transfer technique on an installed network. For spans beyond 50 km, the increased attenuation will require the insertion of optical amplifiers. Although an *a priori* estimation of noise levels is difficult, a higher level of fiber-induced phase noise is expected in the less protected environment of an installed network. This will be counteracted by increasing the phase noise suppression bandwidth closer to the theoretical limit imposed by the round-trip delay.

We thank D. Richardson and P. Petropoulos of the Optoelectronics Research Centre at the University of Southampton for useful discussions and for lending us some of the equipment. This work is funded by the UK Department of Business, Innovation and Skills.

References

1. A. Bauch, J. Achkar, S. Bize, D. Calonico, R. Dach, R. Hlavać, L. Lorini, T. Parker, G. Petit, D. Piester, K. Szymaniec, and P. Uhrich, Metrologia **43**, 109 (2006).
2. L. Maleki, ed., *Proceedings of the 7th Symposium on Frequency Standards and Metrology* (World Scientific, 2009), p. 73.
3. S. M. Foreman, K. W. Holman, D. D. Hudson, D. J. Jones, and J. Ye, Rev. Sci. Instrum. **78**, 021101 (2007).
4. O. Lopez, A. Amy-Klein, M. Lours, Ch. Chardonnet, and G. Santarelli, Appl. Phys. B **98**, 723 (2010).
5. P. A. Williams, W. C. Swann, and N. R. Newbury, J. Opt. Soc. Am. B **25**, 1284 (2008).
6. H. Jiang, F. Kéfélian, S. Crane, O. Lopez, M. Lours, J. Millo, D. Holleville, P. Lemonde, Ch. Chardonnet, A. Amy-Klein, and G. Santarelli, J. Opt. Soc. Am. B **25**, 2029 (2008).
7. F. L. Hong, M. Musha, M. Takamoto, H. Inaba, S. Yanagimachi, A. Takamizawa, K. Watabe, T. Ikegami, M. Imae, Y. Fujii, M. Amemiya, K. Nakagawa, K. Ueda, and H. Katori, Opt. Lett. **34**, 692 (2009).
8. G. Grosche, O. Terra, K. Predehl, R. Holzwarth, B. Lipphardt, F. Vogt, U. Sterr, and H. Schnatz, Opt. Lett. **8**, 531 (2009).
9. K. W. Holman, D. D. Hudson, J. Ye, and D. J. Jones, Opt. Lett. **30**, 1225 (2005).
10. S. T. Dawkins, J. J. McFerran, and A. N. Luiten, IEEE Trans. Ultrason. Ferroelectr. Freq. Control **54**, 918 (2007).
11. B. C. Young, F. C. Cruz, W. M. Itano, and J. C. Bergquist, Phys. Rev. Lett. **82**, 3799 (1999).
12. S. A. Webster, M. Oxborrow, S. Pugla, J. Millo, and P. Gill, Phys. Rev. A **77**, 033847 (2008).
13. E. N. Ivanov, S. A. Diddams, and L. Hollberg, IEEE Trans. Ultrason. Ferroelectr. Freq. Control **52**, 1068 (2005).

High-resolution microwave frequency transfer over an 86-km-long optical fiber network using a mode-locked laser

Giuseppe Marra,[1,2,*] Radan Slavík,[2] Helen S. Margolis,[1] Stephen N. Lea,[1] Periklis Petropoulos,[2] David J. Richardson,[2] and Patrick Gill[1]

[1]*National Physical Laboratory, Hampton Road, Teddington, Middlesex, TW11 0LW, UK*

[2]*Optoelectronics Research Centre, University of Southampton, Southampton, SO17 1BJ, UK*

**Corresponding author: giuseppe.marra@npl.co.uk*

Received October 5, 2010; revised December 15, 2010; accepted January 8, 2011; posted January 14, 2011 (Doc. ID 136149); published February 8, 2011

We demonstrate the transfer of an ultrastable microwave frequency by transmitting a 30-nm-wide optical frequency comb from a mode-locked laser over 86 km of installed optical fiber. The pulse train is returned to the transmitter via the same fiber for compensation of environmentally induced optical path length changes. The fractional transfer stability measured at the remote end reaches 4×10^{-17} after 1600 s, corresponding to a timing jitter of 64 fs. © 2011 Optical Society of America

OCIS codes: 120.3940, 140.4050, 060.2360, 060.5530.

One of the most interesting properties of femtosecond mode-locked lasers is their ability to provide simultaneously a large number of coherent and regularly spaced optical and microwave frequencies. The development of this type of laser has enabled fundamental advances in frequency metrology over the past decade [1]. Over the past few years, they have been identified as ideal candidates for the synchronization of microwave and optical timing signals in particle accelerators and x-ray free-electron lasers [2–4], where sub-100 fs timing jitter is currently required and sub-10 fs is likely to be required for the upcoming generation [5]. They could also find applications in arrays of radiotelescopes, which have very demanding timing requirements for correlation of the signals captured by the different antennas within the array [6]. In both particle physics and radioastronomy facilities, timing signal distribution is typically required over distances ranging from a few hundred meters up to 20 km.

In previous work, we have investigated the use of mode-locked lasers for frequency metrology applications, transferring a high-stability microwave frequency by propagation of an optical frequency comb over 50 km of spooled fiber in the laboratory [7], reporting improved results compared to previous experiments over 3.5 km of optical fiber network [8]. The experiment described in this Letter has three novel aspects. First, we show that similar levels of microwave transfer stability to those demonstrated with spooled fiber [7] can be achieved by propagation of an optical comb over 86 km of an installed fiber network, which is 24 times longer than previously reported [8]. The fiber phase-noise cancellation method is also different, being achieved by returning the pulse train to the transmitter in the same fiber rather than a separate one. Finally, the transfer stability is measured for the first time for time scales longer than 100 s. The stability is demonstrated to reach a few parts in 10^{17} for averaging times above 1000 s, corresponding to a timing jitter of 64 fs.

The experiment was performed on a span of the JANET-Aurora network, a dark-fiber network linking five universities in England. Two fibers are available in the first span from the University of Southampton to a hub located 43 km away. By joining the two fibers at the hub, an 86 km loop with both ends in the laboratory was formed, enabling the measurement of the actual frequency stability delivered to the "user" end. Optical time domain reflectometry measurements revealed that fibers with different loss coefficients per unit length are spliced together in this loop at 33 points. The total optical loss of the 86 km link, measured between the ends available in the laboratory, is approximately 26 dB. The chromatic dispersion was measured to be compatible with that of standard telecommunications fiber and was largely compensated by inserting dispersion-compensating fiber (DCF) modules.

The experimental setup is shown in Fig. 1. The 1.56 μm mode-locked fiber laser emits 150 fs pulses at a repetition rate of 250 MHz. The repetition rate is phase locked to a

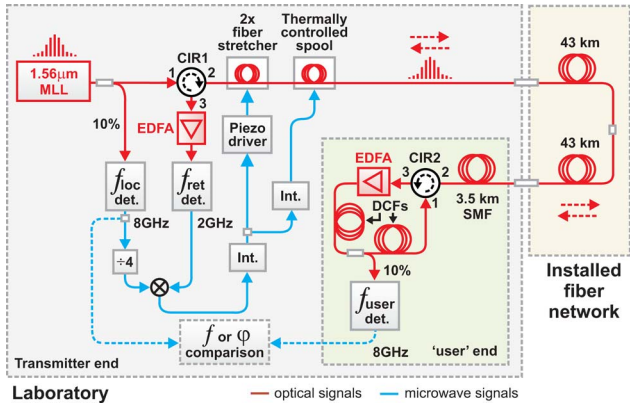

Fig. 1. (Color online) Block diagram of the experimental setup. Each repetition rate detection module consists of a photodiode followed by a tunable bandpass filter and cascaded rf amplifiers. The block marked as f or ϕ comparison is either a phase detector followed by a fast Fourier transform analyzer or an SSB modulator-based frequency counting system. MLL, mode-locked laser; CIR, optical circulator; SMF, single mode fiber; Int., integrator; Det., detection stage; DCF, dispersion-compensating fiber; EDFA, erbium-doped fiber amplifier.

GPS-steered ultrastable quartz crystal oscillator with a feedback bandwidth of a few kilohertz. A 6 m fiber links the fiber laser with the experimental setup on a separate bench. Here, about 20 mW of average optical power is launched into a 90:10 optical splitter. The low-power output provides the "local" reference signal, while the high-power output is used for transmission through the JANET-Aurora network. An optical circulator (CIR1) separates the pulse train launched into the network from that returned after it has travelled a round trip. At the user end of the fiber, another circulator (CIR2) allows a unidirectional erbium-doped fiber amplifier (EDFA) to be used to compensate for the loss of the network. The bandwidth of the EDFA is 30 nm, corresponding to about 15,000 equally spaced optical frequencies being delivered at the user end. The EDFA is followed by two nominally identical DCF modules. DCF1 postcompensates the dispersion of the outward link, ensuring a short pulse duration (less than 100 ps) at the user end. DCF2 precompensates the signal on the return path, such that similarly short pulses are obtained at the transmitter end for the phase comparison. The repetition rate is recovered with a signal-to-noise ratio exceeding 70 dB (in 1 kHz resolution bandwidth) at both the transmitter and user ends. A 3.5 km fiber spool is added between the network and port 2 of CIR2 to better match the dispersion of the DCF modules. The insertion loss of each DCF module is 5.1 dB. We have chosen to position the EDFA before the DCF modules rather than between them to minimize the possibility of nonlinear effects arising from the amplification of short pulses. The optical amplifier increases the average power to 12 dBm of which approximately −3 dBm is delivered to the user photodetector. The remainder is launched back into the same fiber of the network through port 1 of the optical circulator CIR2. At the transmitter end, another optical amplifier at port 3 of CIR1 increases the average power of the returned optical pulse train from −21 to 0 dBm. Fast photodetectors and narrow bandpass filters are used to recover the thirty-second harmonic of the repetition rate (8 GHz) at both the local (f_{loc}) and user end (f_{user}). Because only two fast photodetectors were available, the eighth rather than the thirty-second harmonic of the "returned" repetition rate (f_{ret}) was detected. The phase comparison between the returned and the local (reference) repetition rate is performed at 2 GHz using a microwave mixer preceded by a frequency divide-by-4 prescaler on the f_{loc} arm. The phase-error signal constructed in this way is used to drive two cascaded fiber stretchers and a 300-m-long thermally controlled fiber spool to compensate for the network path length variations, which are primarily due to environmental effects, such as acoustic noise and temperature changes. Time variations of the chromatic-dispersion or polarization-mode dispersion can also degrade the stability of the transferred repetition rate but are secondary contributions compared to the temperature dependence of the index of refraction and the thermal expansion of the fiber.

The total optical path length compensation range of the fiber stretchers is 13 mm and is extended to 50 mm by the thermally controlled spool, allowing a measurement duration of a few hours. A "fast" feedback loop (100 ms time constant) is used to control the fiber stretcher, while a "slow" loop (tens of seconds) controls the thermally controlled spool. The fiber-induced phase noise can be efficiently suppressed at the user end only if the transmitted and returned pulse trains experience the same perturbations over the measurement time scale. This condition is best met when the pulse train propagates in both directions within the same fiber and the optical path changes are slow compared to the round-trip time. Thermal effects on optical devices and short fibers employed in nonbidirectional sections of the optical setup can substantially degrade the long-term stability of the delivered repetition rate. To mitigate this effect, these parts were enclosed in thermally insulating foam.

The optical power launched into the network at the transmitter end is chosen as a compromise between two conflicting requirements. The power needs to be high enough to drive the optical amplifier at the user end with power above the miminum input specification, but it must be low enough to maintain a sufficiently low level of reflections at port 3 of the circulator CIR1, where the repetition rate detected from the returned pulse train and that due to reflections are indistinguishable.

The phase noise of the thirty-second harmonic of the repetition rate delivered to the user is measured against that of the local repetition rate using a microwave mixer. A phase noise reduction of up to 15 dB at low offset frequencies from the carrier is achieved when the phase noise suppression loop is activated (Fig. 2). At 1 Hz offset from the carrier, the single sideband (SSB) phase noise is measured to be −84 dBc/Hz. Integrating the phase noise from 0.1 Hz to 100 kHz, we obtain a timing jitter of 82 fs. However, integrating from 1 to 100 kHz gives a timing jitter of 80 fs, showing that most of the jitter is due to the bumps in the phase noise beyond a few hundred hertz. These arise from self-heterodyne detection of the mode-locked laser noise and the frequencies of the dips correspond to integer multiples of the inverse of the pulse train travel time. These noise bumps could be removed by phase locking a low-noise oscillator to a harmonic of the detected repetition rate with a bandwidth of approximately 1 kHz. We estimate that, in this case, the timing jitter could be reduced to less than 30 fs (1 Hz to

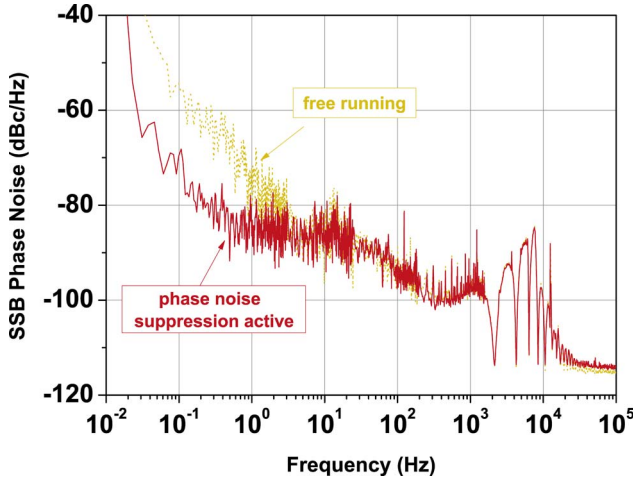

Fig. 2. (Color online) Phase noise of the thirty-second harmonic of the repetition rate measured at the user end when the phase-noise suppression feedback loop is activated and turned off (free running).

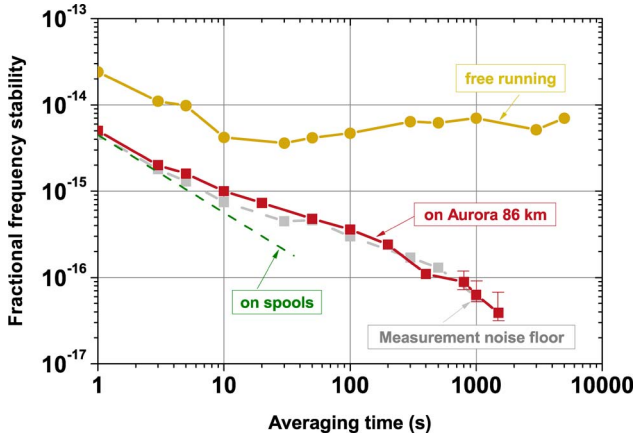

Fig. 3. (Color online) Fractional frequency stability of the thirty-second harmonic of the repetition rate measured at the user end. The stability is limited by the measurement noise floor.

100 kHz). If a better frequency reference was used to stabilize the repetition rate of the mode-locked laser (e.g., if the comb was locked to an optical reference) or if shorter optical links were used, it is likely that similar jitter levels could be achieved without using a phase-locked oscillator.

The feedback control bandwidth was intentionally limited to a few tens of hertz as amplitude noise present at higher frequencies, indistinguishable from phase noise at the output of the error signal phase detector, prevented correct phase-noise cancellation at the user end. The transfer frequency stability, shown in Fig. 3, was measured by a Λ-type counter [9] using a SSB modulator driven by a direct digital synthesizer (DDS), as in our previous experiment [7]. The transfer stability delivered to the user is 5×10^{-15} at 1 s and reaches 4×10^{-17} at 1600 s, corresponding to a timing jitter of 64 fs. The results are limited by the measurement noise floor, which was determined by replacing the fiber network and the DCF modules with optical attenuators of corresponding optical loss. The measurement noise floor was found to be imposed by the noise of the optical amplifiers. We have identified a different EDFA model exhibiting lower noise, but, unfortunately, only one of these amplifiers was available for this experiment. The contribution of the SSB modulator to the measurement noise floor was measured to be 2.2×10^{-16} at 1 s and 1.2×10^{-18} at 1000 s.

The accuracy of the delivered thirty-second harmonic of the repetition rate was measured by comparing the average frequency of the beat between f_{user} and f_{loc} with the frequency produced by the DDS, both averaged over 1600 s. The mean fractional frequency difference was measured to be 1.4×10^{-17}, which is compatible with the frequency stability measured over the same time interval.

In summary, we have demonstrated that a frequency comb can be used for the transfer of a microwave frequency over many tens of kilometers of installed optical network with a stability compatible with state-of-the-art frequency references. Although transfer stabilities five to ten times better have been demonstrated using a microwave transfer technique based on intensity modulation of an optical carrier frequency [10], we expect that comparable performance could be obtained using our method with lower-noise optical amplifiers. The comb transfer technique also has the advantage that a large number of regularly spaced microwave and optical frequencies (rf) can be transferred simultaneously. This is particularly useful in particle accelerators, where both rf devices and mode-locked lasers need to be remotely synchronized with high timing accuracy. The measured timing jitter of the transferred microwave signals is compatible with the timing stability requirements of the current generation of particle accelerators and could be further reduced by adding phase-locked clean-up circuits.

This work is funded by the UK Department for Business, Innovation and Skills. We acknowledge the support of JISC and JANET (UK) for the funding and provision of the JANET-Aurora dark-fiber facility.

References

1. T. Udem, R. Holzwarth, and T. Hänsch, Eur. Phys. J. Special Topics **172**, 69 (2009).
2. J. Kim, J. A. Cox, J. Chen, and F. X. Kärtner, Nat. Photon. **2**, 733 (2008).
3. E. Allaria, C. Callegari, D. Cocco, W. M. Fawley, M. Kiskinova, C. Masciovecchio, and F. Parmigiani, New J. Phys. **12**, 075002 (2010).
4. P. R. Bolton, Int. J. Mod. Phys. B **21**, 527 (2007).
5. F. Loehl, V. Arsov, M. Felber, K. Hacker, B. Lorbeer, F. Ludwig, K. Matthiesen, H. Schlarb, B. Schmidt, A. Winter, S. Schulz, J. Zemella, J. Szewinski, and W. Jalmuzna, in *Proceedings of 11th European Particle Accelerator Conference* (2008), pp. 3360–3362.
6. J. F. Cliche and B. Shillue, IEEE Control Syst. Mag. **26**, 19 (2006).
7. G. Marra, H. S. Margolis, S. N. Lea, and P. Gill, Opt. Lett. **35**, 1025 (2010).
8. K. W. Holman, D. D. Hudson, J. Ye, and D. J. Jones, Opt. Lett. **30**, 1225 (2005).
9. S. T. Dawkins, J. J. McFerran, and A. N. Luiten, IEEE Trans. Ultrason. Ferroelectr. Freq. Control **54**, 918 (2007).
10. O. Lopez, A. Amy-Klein, M. Lours, C. Chardonnet, and G. Santarelli, Appl. Phys. B **98**, 723 (2010).

Dissemination of an optical frequency comb over fiber with 3×10^{-18} fractional accuracy

Giuseppe Marra,[1,2,*] **Helen S. Margolis,**[1] **and David J. Richardson**[2]

[1]*National Physical Laboratory, Hampton Road, Teddington, TW11 0LW, UK*
[2]*Optoelectronics Research Centre, University of Southampton, Highfield Campus, Southampton, SO17 1BJ, UK*

[*]*giuseppe.marra@npl.co.uk*

Abstract: We demonstrate that the structure of an optical frequency comb transferred over several km of fiber can be preserved at a level compatible with the best optical frequency references currently available. Using an optical phase detection technique we measure the noise introduced by the fiber link and suppress it by stabilizing the optical path length. The measured fractional frequency stability of the transferred optical modes is 2×10^{-18} at a few thousand seconds and the mode spacing stability after optical-microwave conversion is better than 4×10^{-17} over the same time scale.

OCIS codes: (120.3940) Metrology; (120.3930) Metrological instrumentation; (140.4050) Mode-locked lasers ; (060.2360) Fiber optics links and subsystems.

References and links

1. M. J. Thorpe, K. D. Moll, R. J. Jones, B. Safdi, and J. Ye, "Broadband cavity ringdown spectroscopy for sensitive and rapid molecular detection," Science **311**, 1595–1599 (2006).
2. S. A. Diddams, L. Hollberg, and V. Mbele, "Molecular fingerprinting with the resolved modes of a femtosecond laser frequency comb," Nature **445**, 627–630 (2007).
3. J. Mandon, G. Guelachvili, and N. Picqué, "Fourier transform spectroscopy with a laser frequency comb," Nature Photon. **3**, 99–102 (2009).
4. F. Krausz and M. Ivanov, "Attosecond physics," Rev. Mod. Phys. **81**, 163–234 (2009).
5. C.-H. Li, A. J. Benedick, P. Fendel, A. G. Glenday, F. X. Kärtner, D. F. Phillips, D. Sasselov, A. Szentgyorgyi, and R. L. Walsworth, "A laser frequency comb that enables radial velocity measurements with a precision of 1 cm s^{-1}," Nature **452**, 610–612 (2008).
6. T. Steinmetz, T. Wilken, C. Araujo-Hauck, R. Holzwarth, T. W. Hänsch, L. Pasquini, A. Manescau, S. D'Odorico, M. T. Murphy, T. Kentischer, W. Schmidt, and T. Udem, "Laser frequency combs for astronomical observations," Science **321**, 1335–1337 (2008).
7. T. Rosenband, D. B. Hume, P. O. Schmidt, C. W. Chou, A. Brusch, L. Lorini, W. H. Oskay, R. E. Drullinger, T. M. Fortier, J. E. Stalnaker, S. A. Diddams, W. C. Swann, N. R. Newbury, W. M. Itano, D. J. Wineland, and J. C. Bergquist, "Frequency ratio of Al^+ and Hg^+ single-ion optical clocks; metrology at the 17^{th} decimal place," Science **319**, 1808–1812 (2008).
8. "STE-QUEST, Space-Time Explorer and QUantum Equivalence Principle Space Test," http://sci.esa.int/ste-quest.
9. C. W. Chou, D. B. Hume, J. C. J. Koelemeij, D. J. Wineland, and T. Rosenband, "Frequency comparison of two high-accuracy Al^+ optical clocks," Phys. Rev. Lett. **104**, 070802 (2010).
10. L.-S. Ma, Z. Bi, A. Bartels, L. Robertsson, M. Zucco, R. S. Windeler, G. Wilpers, C. Oates, L. Hollberg, and S. A. Diddams, "Optical frequency synthesis and comparison with uncertainty at the 10^{-19} level," Science **303**, 1843–1845 (2004).

#157106 - $15.00 USD Received 25 Oct 2011; revised 2 Dec 2011; accepted 10 Dec 2011; published 11 Jan 2012

11. G. Grosche, K. Predehl, S. M. F. Raupach, H. Schnatz, O. Terra, S. Droste, and R. Holzwarth "High performance frequency comparisons over optical fibre," IQEC/CLEO Pacific Rim 2011, Technical Digest CD-ROM (Sydney, Australia, 2011)
12. O. Lopez, A. Amy-Klein, M. Lours, C. Chardonnet, and G. Santarelli, "High-resolution microwave frequency dissemination on an 86-km urban optical link," Appl. Phys. B **98**, 723–727 (2010).
13. G. Grosche, O. Terra, K. Predehl, R. Holzwarth, B. Lipphardt, F. Vogt, U. Sterr, and H. Schnatz, "Optical frequency transfer via 146 km fiber link with 10^{-19} relative accuracy," Opt. Lett. **34**, 2270–2272 (2009).
14. H. Jiang, F. Kéfélian, S. Crane, O. Lopez, M. Lours, J. Millo, D. Holleville, P. Lemonde, C. Chardonnet, A. Amy-Klein, and G. Santarelli, "Long-distance frequency transfer over an urban fiber link using optical phase stabilization," J. Opt. Soc. Am. B **25**, 2029–2035 (2008).
15. P. A. Williams, W. C. Swann, and N. R. Newbury, "High-stability transfer of an optical frequency over long fiber-optic links," J. Opt. Soc. Am. B **25**, 1284–1293 (2008).
16. M. Musha, F.-L. Hong, K. Nakagawa, and K. Ueda, "Coherent optical frequency transfer over 50-km physical distance using a 120-km-long installed telecom fiber network," Opt. Express **16**, 16459–16466 (2008).
17. F. Kéfélian, O. Lopez, H. Jiang, C. Chardonnet, A. Amy-Klein, and G. Santarelli, "High-resolution optical frequency dissemination on a telecommunications network with data traffic," Opt. Lett. **34**, 1573–1575 (2009).
18. S. M. Foreman, K. W. Holman, D. D. Hudson, D. J. Jones, and J. Ye, "Remote transfer of ultrastable frequency references via fiber networks," Rev. Sci. Instrum. **78**, 021101 (2007).
19. K. W. Holman, D. D. Hudson, J. Ye, and D. J. Jones, "Remote transfer of a high-stability and ultralow-jitter timing signal," Opt. Lett. **30**, 1225–1227 (2005).
20. G. Marra, R. Slavík, H. S. Margolis, S. N. Lea, P. Petropoulos, D. J. Richardson, and P. Gill, "High-resolution microwave frequency transfer over an 86-km-long optical fiber network using a mode-locked laser," Opt. Lett. **36**, 511–513 (2011).
21. C.-W. Kang, T.-A. Liu, R.-H. Shu, C.-L. Pan, and J.-L. Peng, "Simultaneously transfer microwave and optical frequency through fiber using mode-locked fiber laser," in *Conference on Lasers and Electro-Optics/International Quantum Electronics Conference*, OSA Technical Digest, (Optical Society of America, 2009), p. JWA70.
22. Y.-F. Chen, J. Jiang, and D. J. Jones, "Remote distribution of a mode-locked pulse train with sub 40-as jitter," Opt. Express **14**, 12134–12144 (2006).
23. E. Ivanov, S. Diddams, and L. Hollberg, "Study of the excess noise associated with demodulation of ultra-short infrared pulses," IEEE Trans. Ultrason., Ferroelectr. Freq. Control **52**, 1068–1074 (2005).

1. Introduction

A decade after the optical frequency comb technique transformed the field of frequency metrology, these devices have already found applications in other science areas as diverse as spectroscopy [1–3], attosecond physics [4] and astrophysics [5, 6]. When combined with atomic references, tests of fundamental physics that would have been unthinkable only a few years ago, and whose outcome could open a whole new era for physics, can now be performed [7, 8]. With optical frequency standards currently exhibiting a fractional accuracy better than 10^{-17} [9] and optical frequency combs making this accuracy available across a wide spectrum [10], new experiments could be devised in a wide range of research fields if ultra accurate optical frequencies were to be made available beyond the walls of metrology laboratories.

Much research has recently been undertaken on the transfer of ultra-stable frequencies over long lengths of optical fiber, driven by the need to compare state-of-the-art optical clocks in distantly located laboratories. Steps towards a world-wide optical fiber network linking these laboratories have already been undertaken, with a 900 km link now operative in Germany [11] and shorter links being progressively set up across Europe. Remarkable results have been demonstrated for transfer of a single microwave or optical frequency over long spans of dark fiber, using phase noise cancellation techniques to compensate for the noise introduced by environmental perturbations to the fiber [12–16]. Tests have also been performed on fiber carrying internet traffic [17], since international clock comparisons will inevitably need to use commercial optical networks. However in dedicated research networks a large optical bandwidth is available, making it possible to transfer a frequency comb rather than a single optical frequency [18]. It has been shown that microwave frequencies can be transferred over research networks with a stability and accuracy compatible with state-of-the-art frequency standards by using the mode spacing of an optical frequency comb [19, 20]. However, although preliminary experi-

#157106 - $15.00 USD Received 25 Oct 2011; revised 2 Dec 2011; accepted 10 Dec 2011; published 11 Jan 2012

ments were carried out on a 1.5 km uncompensated fiber link [21], the issue of how accurately it is possible to preserve the optical comb structure over a long length of fiber has not been addressed until now. Here, we show for the first time that this is possible at a level exceeding that of the best frequency standards available today. We demonstrate that the mode frequencies of an optical comb can be transferred over a several km fiber link with a fractional accuracy better than 3×10^{-18}, a factor of three more accurate than the best optical clock reported to date [9]. We also test the stability of the mode spacing after optical-microwave conversion and find that it is preserved to better than 2×10^{-15} at 1 s and 4×10^{-17} for averaging times greater than 1000 s. This performance is achieved by using an optical phase detection technique to cancel the environmentally induced fiber noise, an approach that provides orders of magnitude higher sensitivity than techniques based on microwave detection [18].

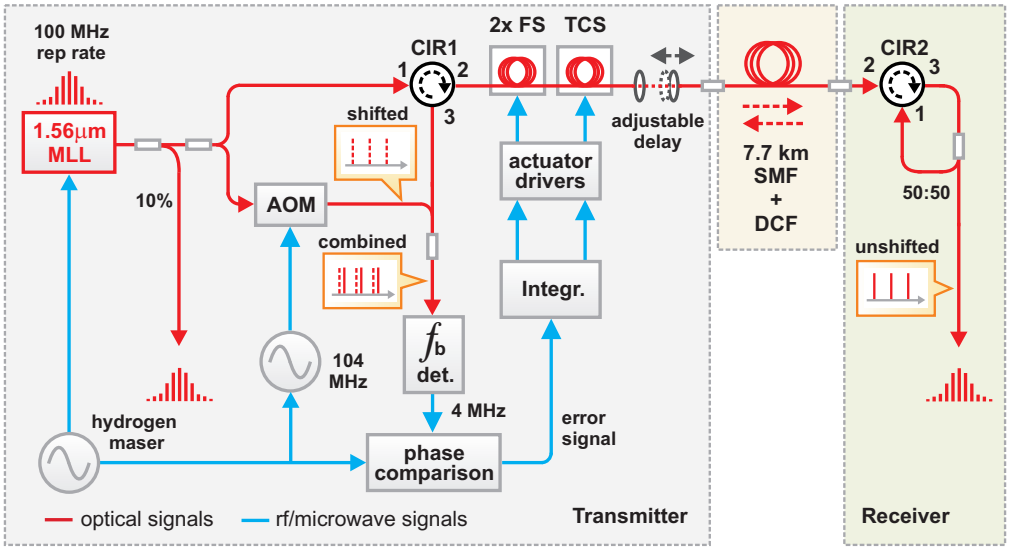

Fig. 1. Experimental setup. The detection stage (f_b det.) consists of a photodiode, a filter and cascaded amplifiers. MLL: mode-locked laser; CIR: circulator; FS: fiber stretchers; TCS: thermally controlled spool; Integr.: integrator; AOM: acousto-optic modulator; SMF: single-mode fiber; DCF: dispersion compensating fiber.

2. Fiber noise suppression principle

The fiber phase noise cancellation principle, which is related to one previously used for the transfer of ultra-short pulses over 60 m of fiber [22], is illustrated in Fig. 1. A commercial 1.56 μm amplified erbium-doped mode-locked fiber laser (MLL) generates sub-150 fs optical pulses at a repetition rate $f_r = 100$ MHz. In the frequency domain this pulse train corresponds to a frequency comb with approximately 10^5 optical modes extending over an optical bandwidth of approximately 100 nm. The frequency of each optical mode can be described by $f_m = m f_r + f_0$ where m is an integer, f_r is the repetition rate and f_0 is the carrier-envelope offset frequency. Both f_r and f_0 are stabilized to a 10 MHz signal from a hydrogen maser. Approximately 10 m of SMF-28 fiber are used to connect the experimental setup to the laser source, which is located in a different part of the laboratory, broadening the pulse duration to approximately 17 ps before it enters the first power splitter. The 90% output of this 90:10 splitter is used to propagate the comb over 7.7 km of spooled single mode fiber (SMF) to the receiver end where a portion

#157106 - $15.00 USD Received 25 Oct 2011; revised 2 Dec 2011; accepted 10 Dec 2011; published 11 Jan 2012

is returned to the transmitter end via the same fiber. Both ends of the fiber are located in the same laboratory so that the accuracy and stability of the signal that has travelled 7.7 km can be compared to the signal injected at the input of the fiber. The forward and backward travelling pulse trains from the mode-locked laser are separated using optical circulators (CIR1, CIR2) and a dispersion compensating fiber module (DCF) recompresses the pulses to a duration of less than 100 ps. A free-space delay line is adjusted to achieve appropriate temporal overlap between the local (10% output of the power splitter) and the returned pulse trains. The returned frequency comb is combined with the original comb after the latter has been frequency shifted by f_{AOM} = 104 MHz using an acousto-optic modulator (AOM) and the beat notes between their optical modes are detected with a photodiode. The total length of out-of-loop fiber in the measurement setup was a few meters, with the lengths travelled by the shifted and unshifted combs being similar in order to improve the common-mode rejection of environmental effects that could increase the measurement noise floor.

The Doppler-induced frequency shift on the optical comb mode $f_p = p f_r + f_0$, due to environmental perturbations on the fiber, can be written as

$$\delta f_p = \frac{1}{c}\frac{d[n(f_p)l]}{dt} f_p \tag{1}$$

where $n(f_p)$ and $d[n(f_p)l]/dt$ are the refractive index of the fiber and the instantaneous rate of change of the optical path length at frequency f_p, and c is the speed of light in vacuum. This frequency shift can be measured by detecting any of the beat frequencies $f_{\mathrm{b}p,m}$ arising from interference between the optical modes f_m and f_p of the AOM-shifted and unshifted combs respectively:

$$f_{\mathrm{b}p,m} = f_{\mathrm{AOM}} - (f_p - f_m) - \frac{1}{c}\frac{d[n(f_p)l]}{dt} f_p \tag{2}$$

where we have assumed that the optical frequencies generated by the mode-locked laser are constant over the round trip time so that no noise is detected due to the self-heterodyne effect. This assumption is reasonable when the self-heterodyne noise is lower than the phase noise introduced by enviromental perturbations to the fiber link, which is the case in our experiment for frequencies within the feedback bandwidth of approximately 2 kHz.

These detected beats give rise to a current from the photodetector that can be described as

$$i(t) \propto \sum_{m=m_1}^{m_2} \sum_{p=p_1}^{p_2} \cos 2\pi \left[f_{\mathrm{AOM}} - (p-m) f_r - \frac{1}{c}\frac{d[n(f_p)l]}{dt} f_p \right] t \tag{3}$$

where m_1, m_2, p_1 and p_2 define the range of comb modes that contribute to the signal. Equation (3) illustrates the gain in sensitivity achieved using optical rather than microwave phase detection techniques; the frequency shift due to changes in the optical path length is greater by a factor $f_p/f_r \sim 10^6$. In our experiment f_{AOM} was 104 MHz, f_r was 100 MHz and we chose to detect the lowest frequency beat such that $p = m+1$, in which case

$$i(t) \propto \sum_{p=p_1}^{p_2} \cos 2\pi \left[f_{\mathrm{AOM}} - f_r - \frac{1}{c}\frac{d[n(f_p)l]}{dt} f_p \right] t. \tag{4}$$

For simplicity, Eqs. (3) and (4) assume that all comb modes have the same power and interfere with the same phase offset. In reality, the residual dispersion of the fiber link means that destructive interference between the beat frequencies will reduce $i(t)$, reducing the signal-to-noise ratio of the extracted beat signal. However, this effect will be partially counteracted by variations in mode power across the spectrum, and in practice a sufficiently high signal is detected.

#157106 - $15.00 USD Received 25 Oct 2011; revised 2 Dec 2011; accepted 10 Dec 2011; published 11 Jan 2012

The lowest frequency beat, 4 MHz in our experiment, is amplified and phase compared with a maser-referenced synthesizer, generating an error signal which, after integration, is applied to two fiber stretchers and a thermally controlled fiber spool to compensate for fast (up to a few kHz) and slow phase fluctuations respectively.

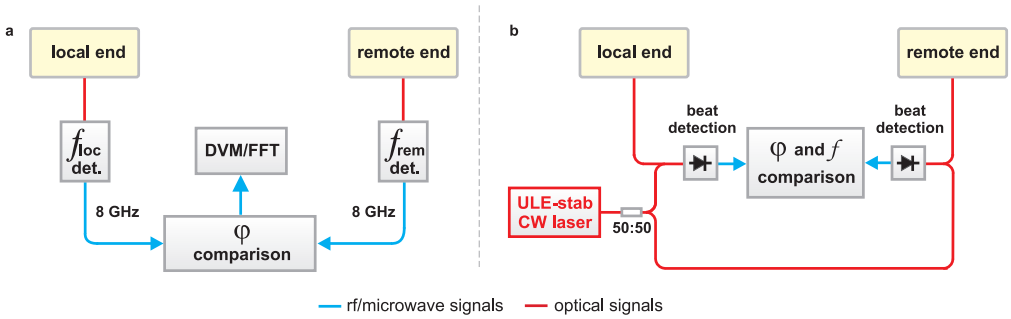

Fig. 2. Experimental setups for phase noise and frequency stability measurement of the mode spacing (a) and transferred optical modes (b). The detection stages f_{loc} and f_{rem} consist of a fast photodiode, a narrow bandpass filter and microwave amplifiers.

3. Results

To test the precision with which the full structure of the optical frequency comb can be preserved when transmitted over several-km-scale fiber links, two comb parameters must be measured: the repetition rate f_r (mode spacing) and the frequency of a selected optical mode.

The measurement of the mode spacing stability, accuracy and phase noise is performed by phase comparing the 80th harmonic (8 GHz) of f_r at the receiver end of the fiber with that detected directly at the output of the laser using a microwave mixer (Fig. 2a). The power spectral density of the phase noise fluctuations between 0.1 Hz and 100 kHz was measured with an FFT analyser. When the noise cancellation is activated, the measured phase noise is -91 dBc/Hz at 1 Hz offset from the carrier (Fig. 3a), very close to the noise measured when the 7.7 km SMF and the DCF are replaced by a 2 m fiber and an attenuator set to provide the same overall loss. We observe suppression of both the thermally induced fiber noise at low offset frequencies (by up to 20 dB) and the acoustic noise around a few hundred hertz. By integrating the phase noise between 1 Hz and 100 kHz we calculate the timing jitter to be less than 17.5 fs. The measurement of the frequency stability was achieved by converting the output voltage of the microwave mixer, logged every 0.5 s with a digital voltmeter (measurement bandwidth 7 Hz), into phase changes. The measured stability is shown in Fig. 3b. The Allan deviation of the data taken within the first two hours is 1.8×10^{-15} at 1 s and 7×10^{-17} at 100 s. Over longer periods, the signal-to-noise ratio (SNR) of the error signal degrades due to changes in the polarization of the returned optical signal with respect to that of the local signal. When the SNR falls below approximately 25 dB (in a 10 kHz bandwidth), we notice a degradation of the fractional frequency stability to 2.2×10^{-15} at 1 s and 1×10^{-16} at 100 s. This problem could be alleviated by using an optical amplifier to increase the power of the returned optical signal in order to maintain the SNR above this level. An alternative approach could be to insert an automatic polarization control system on the returned pulse train. However, even with the SNR degradation the fractional transfer stability is better than 4×10^{-17} at 1000 s, corresponding to a timing jitter of 40 fs.

To determine the stability and accuracy of the transferred optical modes a continuous-wave

#157106 - $15.00 USD Received 25 Oct 2011; revised 2 Dec 2011; accepted 10 Dec 2011; published 11 Jan 2012

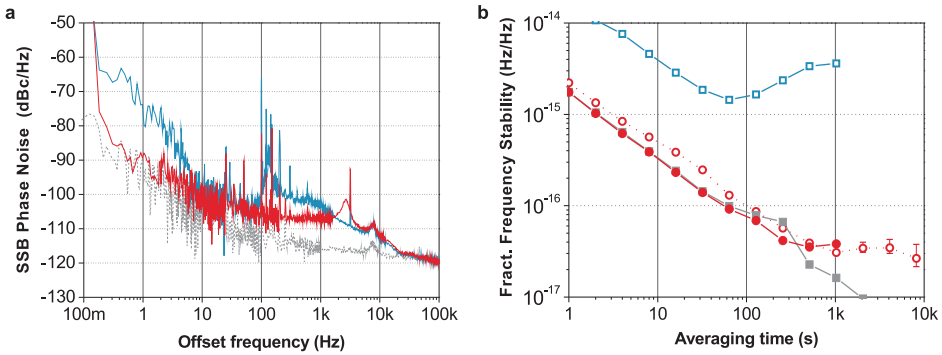

Fig. 3. Single sideband phase noise (a) and frequency stability (b) of the 80th harmonic (8 GHz) of the fundamental frequency spacing measured at the remote end. a) blue: fiber noise not suppressed; red: fiber noise suppressed; grey: measurement noise floor. b) Open blue squares: fiber noise not suppressed; open red circles: fiber noise suppressed; filled red circles: fiber noise suppressed, first two hours; filled grey squares: measurement noise floor.

(CW) 1542 nm laser stabilized to an ultra-low-expansion glass optical cavity is used as a common reference against which a selected optical comb mode frequency is measured before and after the fiber link (Fig. 2b). At each end two 35 MHz beats are generated between a selected comb line and the CW laser. At the remote end, in order to achieve a suitably high SNR of the beat, an erbium-doped optical amplifier is used at port 3 of CIR2. A frequency comb extending over 30 nm and with an average power of 2 mW is available at the output of the amplifier. The broadband noise of the 35 MHz beats is filtered using tracking oscillators with a bandwidth of approximately 200 kHz. Any difference observed between the two beat frequencies arises from the fiber noise since changes in the frequency of the CW laser are common mode. The lengths of the fibers linking the local end, the remote end and the CW laser to the photodetectors were approximately 1 m in each case. The power spectral density of the phase fluctuations between the two beat frequencies (and hence of the transferred optical mode) is measured with a digital phase detector followed by an FFT analyser. The digital phase detector has a linear range extending over 256π, sufficient to detect correctly the fast and slow phase fluctuations. As shown in Fig. 4a, the fiber-induced phase noise is reduced by up to 55 dB when the phase noise suppression loop is activated and is -40 dBc/Hz at 1 Hz offset from the carrier. Above 10 kHz, bumps in the phase noise are visible due to self-heterodyne detection of the mode-locked laser noise. The timing jitter measured from 1 Hz to 100 kHz is 5.2 fs corresponding to a phase change of approximately 2π. The transfer stability is measured in two ways: by phase comparing the two beats with the digital phase detector and by synchronously counting their frequencies. The difference between the two methods is the measurement bandwidth (7 Hz for phase comparison and approximately 200 kHz for frequency counting). The frequency stability calculated from the phase data is shown in Fig. 4b and is 4×10^{-17} at 1 s and approximately 2×10^{-18} for timescales of a few thousand seconds (corresponding to a timing jitter smaller than 10 fs). The mean frequency offset between the two beats corresponds to a transfer accuracy for the optical mode frequency of 2.6×10^{-18}. When a 200 kHz measurement bandwidth is used, the frequency stability is 5×10^{-15} at 1 s and reaches a few parts in 10^{18} at a few thousand seconds. The calculated accuracy in this case is 2.9×10^{-18}. We achieve these results despite the fact that we use only a single actuator to compensate for changes in both f_0 and f_r. Even better performance could potentially be achieved if a more sophisticated

#157106 - $15.00 USD Received 25 Oct 2011; revised 2 Dec 2011; accepted 10 Dec 2011; published 11 Jan 2012

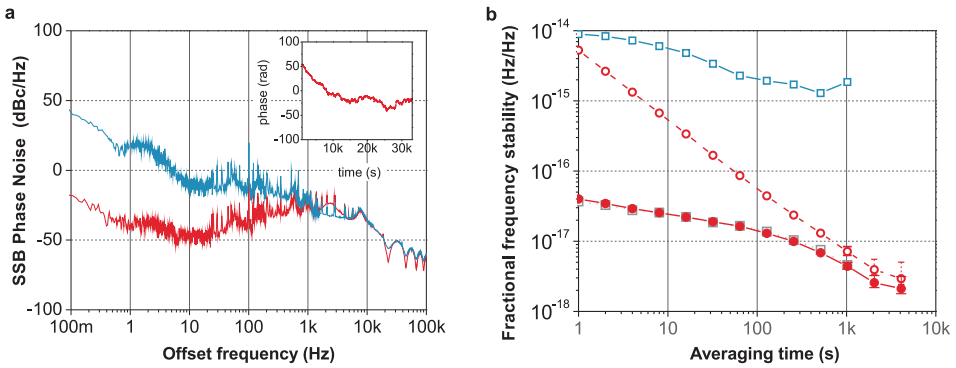

Fig. 4. Phase noise (a) and frequency stability (b) of the sample optical mode at the remote end. a) suppression inactive (blue); active (red). Inset: phase evolution over time of the optical mode when the suppression is active. b) Without (open blue squares) and with noise suppression, 7 Hz (filled red circles) and 200 kHz (open red circles) measurement bandwidth; open grey squares: measurement noise floor.

control system, such as that proposed in [18], were used to compensate for changes in group delay dispersion. Although the accuracy of the transferred repetition rate is worse than that of the optical modes because of noise introduced in optical-to-microwave conversion [23], it is still sufficient to demonstrate that the stability and accuracy measured for a single optical mode is preserved over the transferred optical bandwidth BW_{opt} (approximately 4 THz) since $(\delta f_r/f_r \times BW_{opt})1/f_{opt} \approx 8 \times 10^{-19}$. We also note that the actual stability delivered to a user at the remote end would actually be better than we measured with the frequency counter as no self-heterodyne noise would be present.

4. Conclusion

These results demonstrate that it is possible to preserve the overall structure of an optical frequency comb when it is transferred over several-km lengths of optical fiber, at a level better than the best optical frequency standards currently available. We measure the stability of the optical frequency of a sample mode of the comb to be comparable to that achieved when only a single optical frequency is transferred [13–17] rather than the tens of thousands transferred in our experiment. This frequency comb transfer technique could potentially be used over much longer distances (>100 km) by adding optical amplifiers, possibly bidirectional, to compensate for the extra loss of the fiber. We note that, employing a different phase detection technique, we have previously transferred a microwave frequency over 86 km of installed fiber using a frequency comb [20]. Microwave stability similar to that demonstrated in this paper was achieved in that earlier work, although the stability of optical modes was not measured. As for other frequency transfer techniques, the performance could be degraded by the longer round trip time which limits the fiber noise cancellation bandwidth. However, the degree of degradation will depend on the measurement bandwidth required by the user. For a similar measurement bandwidth to that used in this experiment (7 Hz), transfer over many hundreds of km would be possible.

The results make it possible to envisage the dissemination of highly stable and accurate optical frequency combs from metrology laboratories to other remotely located users. Simultaneous distribution to multiple users can be implemented easily and at low cost since only one AOM is required regardless of the number of fiber links. The availability of a wide comb of

#157106 - $15.00 USD Received 25 Oct 2011; revised 2 Dec 2011; accepted 10 Dec 2011; published 11 Jan 2012

optical frequencies with stability and accuracy matching that of state-of-the-art optical clocks and, equivalently, an optical pulse train with ultra-low timing jitter could enable researchers from different scientific areas to devise new experiments and further extend the applications of optical frequency combs in the years to come.

Acknowledgments

This work was funded by the UK National Measurement System as part of the Pathfinder Metrology programme. We thank Radan Slavík from the University of Southampton for lending us some of the equipment for this experiment and for useful discussions. DJR acknowledges partial support for his contribution to this research from the UK Engineering and Physical Sciences Research Council through grant number EP/I01196X/1.

#157106 - $15.00 USD Received 25 Oct 2011; revised 2 Dec 2011; accepted 10 Dec 2011; published 11 Jan 2012