0.1 Fine time generation and distribution

0.1.1 Background

The goal of a timing system is to provide a common notion of time in a distributed environment. This notion of time is usually the result of counting ticks of a clock signal from an arbitrary instant. The clock signal is ideally of perfect periodicity and stability. Real-world clocks, however, present imperfections [1] in both amplitude and phase as expressed in eq. 1.

$$a(t) = A(1+\alpha(t))\sin(\omega t + \varphi(t)) \tag{1}$$

Amplitude noise can often be controlled through hard amplitude limiters or automatic gain stages. These signals then do not suffer from amplitude modulation, so we will ignore the $\alpha(t)$ term from now on. The random variations in the zero-crossing of the pseudo-periodic signals arise from the $\varphi(t)$ term, usually called phase noise. Ignoring amplitude modulation, eq. 1 can be re-written as

$$a(t) = A\sin\left(\omega\left(t + \frac{\varphi(t)}{\omega}\right)\right) \tag{2}$$

showing that the $\frac{\varphi(t)}{\omega}$ term, which has dimensions of time, represents the time deviations in zero-crossing between the perfect and the imperfect periodic waveforms. $\varphi(t)$ is a random signal whose rms value is in principle a good indicator of clock quality. Dividing that rms value by ω gives the clock jitter. Phase noise is important because this imperfect clock is typically distributed to many receivers, where local counting is done and the common notion of time is generated. In order to compensate for delays in cables and fibers, a constant correction is applied to the local time base, but this assumes the clock is a perfect copy of itself T seconds ago, where T is any multiple of the clock period. If this is not the case, as in all real-life clocks, the delay compensation mechanism does not fully achieve its goal.

0.1.1.1 Phase noise and jitter

Unfortunately, all clocks ultimately diverge in phase and even frequency, in such a way that the rms calculation of jitter gets bigger and bigger as the averaging time grows. In order to tackle this problem, it is useful to work in the frequency domain. The Fourier transform of $\varphi(t)$, noted $\Phi(f)$ has the same energy as the time-domain signal. This result, expressed mathematically in eq. 3, is known as Parseval's theorem [2]:

$$\int_{-\infty}^{+\infty} |\varphi(t)|^2\, dt = \int_{-\infty}^{+\infty} |\Phi(f)|^2\, df \tag{3}$$

The units of the left-hand side (LHS) of eq. 3 are $rad^2 \cdot s$. A real-life signal would be bounded in time. If we call $\varphi_T(t)$ a signal which is non-zero only between times $\frac{-T}{2}$ and $\frac{T}{2}$, its Fourier transform is:

$$\Phi_T(f) = \int_{-T/2}^{+T/2} \varphi_T(t) e^{-j2\pi f t}\, dt \tag{4}$$

Re-writing eq. 3 with the truncated signal and diviving both sides by T we have:

$$\frac{1}{T}\int_{-T/2}^{+T/2} |\varphi_T(t)|^2\, dt = \int_{-\infty}^{+\infty} \frac{|\Phi_T(f)|^2}{T}\, df \tag{5}$$

Since the LHS of eq. 5 is clearly a measure of the power of the signal, the term $\frac{|\Phi_T(f)|^2}{T}$ in the RHS can be interpreted as a Power Spectral Density (PSD). In fact, the Wiener-Khintchine theorem [3] tells us that

$$S_\varphi^{II}(f) = \lim_{T\to\infty} \frac{1}{T} |\Phi_T(f)|^2 \tag{6}$$

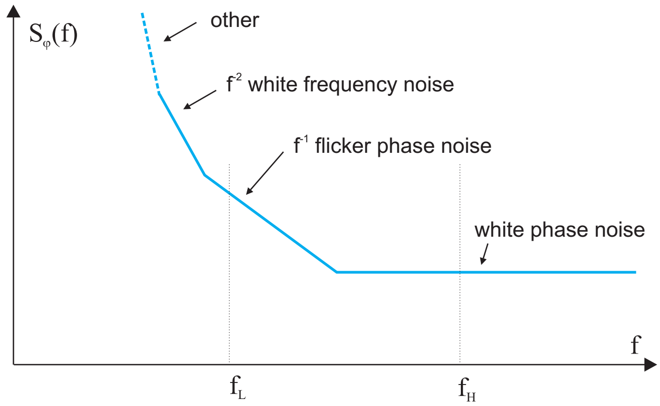

Fig. 1: One-sided PSD of phase noise for a typical oscillator.

where $S_\varphi^{II}(f)$ is the two-sided PSD of the random process $\varphi(t)$. Multiplying by two, we get the one-sided PSD which is the most usual measure of oscillator phase noise. It is also customary to average m finite-time measurements to get an approximation of the one-sided PSD:

$$S_\varphi(f) \approx \frac{2}{T} \left\langle |\Phi_T(f)|^2 \right\rangle_m \tag{7}$$

Taking the square root of eq. 5 we would have the phase noise rms value, and dividing the result by the nominal frequency gives the jitter. The problem, as we said, is that increasing the integration limits results in bigger and bigger measured jitter.

In real life, however, an application is only sensitive to jitter generated between two finite limits in the PSD curve. Figure 1 shows a typical plot of one-sided PSD ($S_\varphi(f)$) of the phase noise for an oscillator. Integration limits are set between f_L and f_H. Phase noise below f_L corresponds to variations which are so slow as to be common mode for all timing receivers under all circumstances. For example, for a machine with a repetition rate of 50 Hz such as CLIC, phase noise below say 1 mHz will give an almost constant contribution during the 20 ms cycle and therefore will not affect the performance of the timing system. In addition, it is estimated that all perturbations below 5 Hz in Fourier frequency can be dealt with by appropriate inter-pulse feedback strategies. Reasons for establishing an upper limit in integration stem mainly from the inability of some systems to react to such fast variations, i.e. to limitations in bandwidth. These limitations can be in electronics, such as the bandwidth of the input stage of a digital gate, or in electro-mechanical systems such as an RF accelerating cavity. It is important to justify lower and upper integration limits for a given application based on both requirements and an intimate knowledge of the system.

Figure 1 also illustrates different types of noise, which can be identified by the different slopes of their PSDs in a log-log graph [1]. White phase noise dominates the high frequency area and has a flat distribution. Moving towards lower frequencies, we find flicker (pink) phase noise, which is characterized by a PSD scaling as f^{-1}. Since frequency is the derivative of phase, white frequency noise – arising from white noise in the frequency-setting elements of an oscillator – features an f^{-2} slope in the phase noise PSD diagram. Higher order f^{-n} terms can also be present. This low-frequency area of the graph will feature quick divergence under integration, and corresponds to the problematic long time-spans mentioned earlier for the time-domain representation.

0.1.1.2 Phase-locked loops

Phase-locked loops [4] are an invaluable tool in cleaning up the jitter of clocks, among many other possible applications. Figure 2 depicts their internal structure.

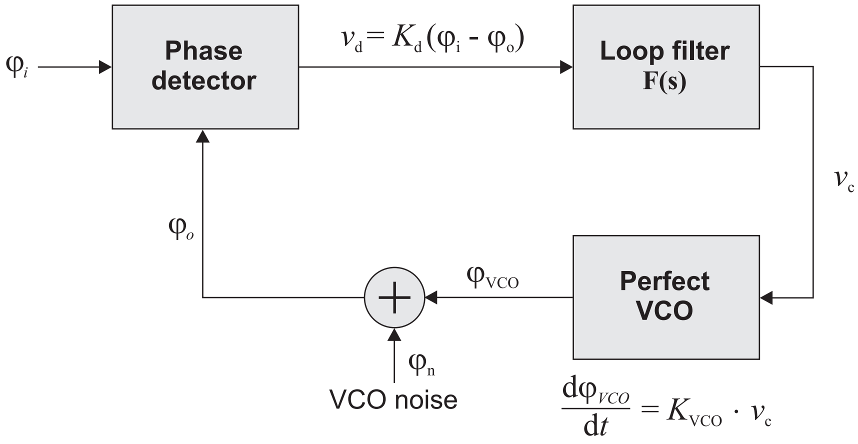

Fig. 2: Block diagram of a phase-locked loop.

The phase detector (PD) block generates an output voltage v_d proportional to the phase difference between the input and output of the PLL. In Laplace space, its output is therefore

$$V_d(s) = K_d\left(\Phi_i(s) - \Phi_o(s)\right) \tag{8}$$

The next block after the phase detector is the loop filter, which outputs the control signal for the Voltage-Controlled Oscillator (VCO):

$$V_c(s) = F(s) \cdot V_d(s) \tag{9}$$

The VCO outputs a signal with a frequency proportional to its input voltage. Since frequency is the derivative of phase, this means that the phase of the signal at the output of the VCO is proportional to the integral of the VCO control voltage:

$$\Phi_{VCO}(s) = \frac{K_{VCO} \cdot V_c(s)}{s} \tag{10}$$

Since there are no perfect VCOs, we have included a VCO noise source in the diagram, contributing phase φ_n. Calculating the output phase φ_o from the two sources in the diagram (reference input phase φ_i and VCO phase noise φ_n) again in Laplace space gives

$$\Phi_o(s) = H(s) \cdot \Phi_i(s) + E(s) \cdot \Phi_n(s) \tag{11}$$

where $H(s)$ is called the system transfer function, defined as

$$H(s) = \frac{K_{VCO} K_d F(s)}{s + K_{VCO} K_d F(s)} \tag{12}$$

and $E(s)$ is the so called error transfer function, defined as

$$E(s) = 1 - H(s) = \frac{s}{s + K_{VCO} K_d F(s)} \tag{13}$$

In typical clock-cleaning applications, $H(s)$ is a low-pass filter, while $E(s)$ is high-pass. Cut-off frequencies are dictated by PLL parameters, and most importantly the loop filter $F(s)$. The PSD of the phase noise of φ_i will be filtered by $|H(s)|^2$ while the phase noise PSD of the VCO will be filtered by $|E(s)|^2$. This means that the low frequency noise in the PSD of φ_o will come from the reference φ_i and the high-frequency noise will come from φ_n. The transition from one noise source to the other will be at

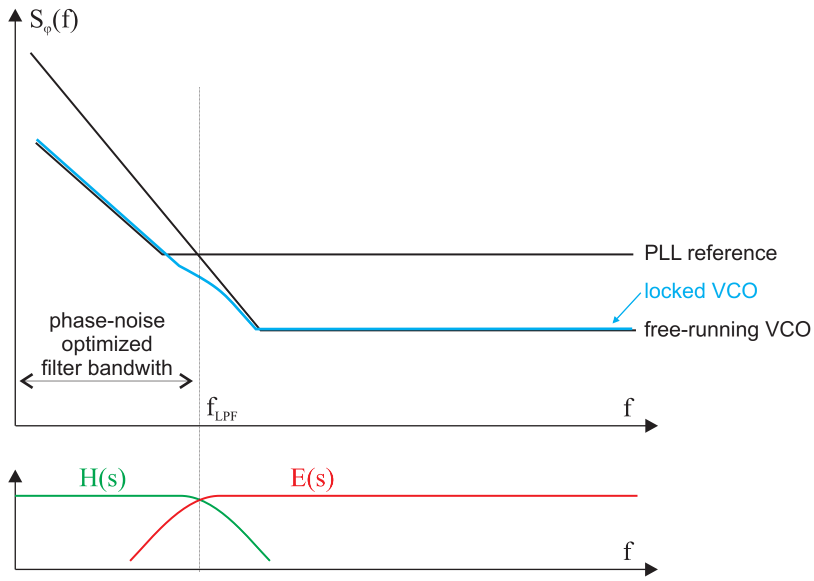

Fig. 3: Optimal choice of PLL bandwidth for jitter-cleaning applications.

a frequency determined by the loop parameters. After careful study of the PSDs of φ_i and φ_n it is the task of the designer to choose a cut-off frequency that will minimize overall area under the φ_o PSD curve, and consequently time-domain jitter. In typical systems – like the transmission of a very stable clock over a channel which adds high-frequency noise – the VCO is worse than the reference at low frequencies and better at high frequencies. The point in frequency where the two PSD plots (reference and VCO) cross is in that case an optimum setting for PLL bandwidth, as shown in figure 3.

In figure 2 the phase detector is shown as perfect, with no noise added to it as for the VCO. In practice, phase detector noise is also a concern, but mathematically it is equivalent to reference noise, so that the above formalism can be applied, replacing reference noise by reference + phase detector noise.

0.1.2 CLIC timing requirements

The main requirements in terms of clock signal distribution come from three subsystems: the Low Level RF control for the drive beam, instrumentation for the main beam, and the two-beam acceleration scheme. The following paragraphs examine each one separately. The conclusion is that CLIC will need the distribution of precise clock signals at the level of some tens of femtoseconds of jitter over different bandwidths.

0.1.2.1 Drive beam RF system requirements

In [5] the jitter of the 1 GHz field in the accelerating cavities of the drive beam is specified as 50 fs integrating between 5 kHz and 20 MHz. It is also said that with appropriate feed-forward control in the main beam, this figure could be relaxed by a factor of 10. However, the reference phase noise fed to the LLRF system is only responsible for a small percentage of the final jitter in the electromagnetic field. Taking this contribution to be 10% results in a specification of 50 fs for the jitter of the reference clock signal distribution to each one of the 326 accelerating structures in each Linac.

0.1.2.2 Beam instrumentation requirements

Longitudinal profile monitors could be based on distributed lasers and changes in optical properties of birefringent materials induced by the main beam. These monitors have the mission of measuring 150 fs bunches with a resolution of 20 fs [6]. The precision required from the clock signal allowing

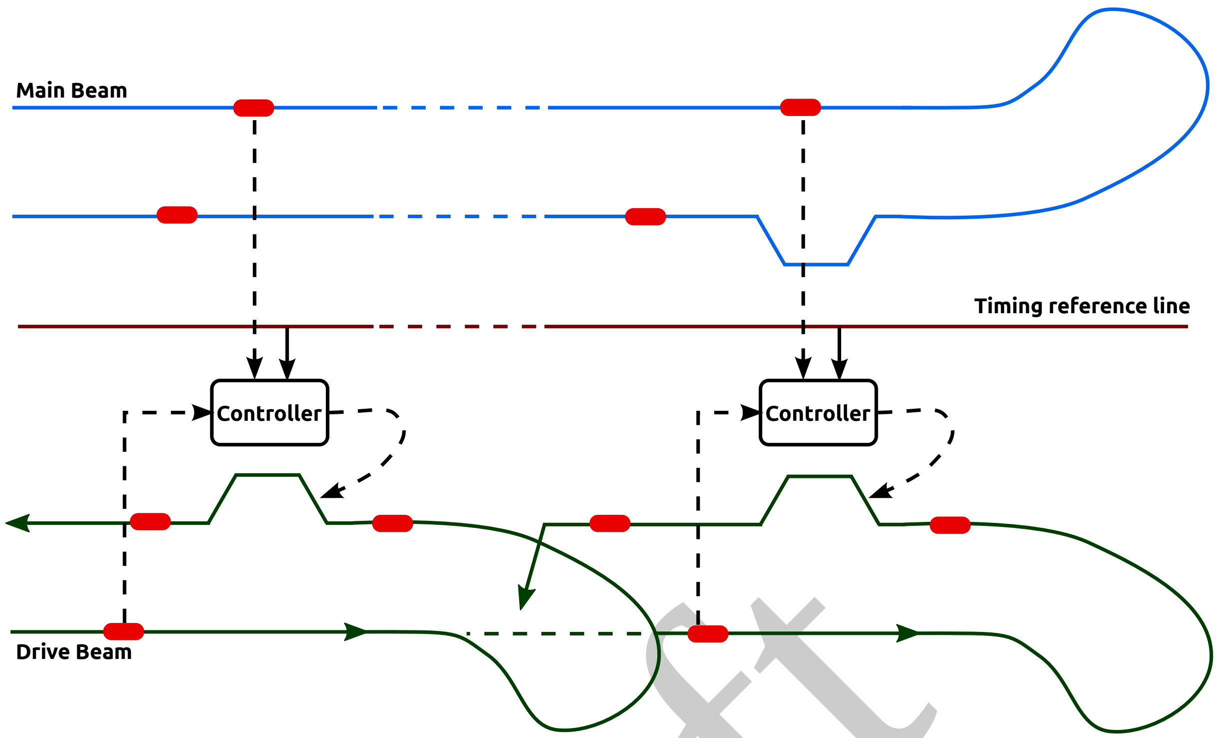

Fig. 4: Timing reference line usage for inter-beam synchronization (adapted from [7]).

to synchronize the lasers with the beam would therefore be in the few tens of femtoseconds, integrated between 5 Hz and a few hundreds of kHz (the bandwidth of the PLL locking the laser to the reference).

0.1.2.3 Two-beam acceleration system requirements

The two-beam acceleration scheme in CLIC requires a very precise synchronization between the drive beam and the main beam. Figure 4 depicts a possible synchronization solution. A controller measures the phases of the drive and main beams with respect to a reference line, and uses that information to control the amplitude of a kicker pulse which modifies the trajectory of the drive beam in order to keep it well synchronized with the main beam. The required precision of this alignment is around 40 fs [7], so the timing reference precision clearly needs to be better than that. The bandwidth of the kickers, in the several MHz region, would set a natural upper limit for integration of phase noise.

0.1.3 Technical description

This subsection describes the two main types of solutions currently developed and tested for the distribution of precise timing signals in the femtosecond realm [8]. Both types have achieved synchronizations better than 20 fs over distances of several hundreds of meters.

0.1.3.1 Continuous Wave systems

Figure 5 shows a simplified view of a system based on a Continuous Wave (CW) laser modulated in amplitude by the RF or microwave signal to transmit. A fraction of the ligth reaching the destination bounces back from a Faraday Rotator Mirror (FRM) and is in the process shifted in frequency by 100 MHz using a Frequency Shifter (FS). Another FRM bounces off a sample of the light signal as it gets out of the source. These two signals, upon mixing together, produce a beat at 100 MHz whose phase can be easily measured. The key feature of this system is that the heterodyning process preserves phases, so a phase shift induced in the optical frequency by a change of length in fiber will show up as exactly the same

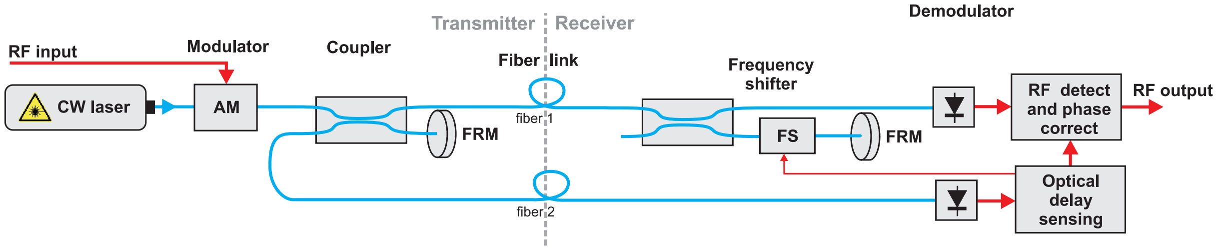

Fig. 5: Continuous wave system for femtosecond timing distribution (simplified).

phase shift at the RF (100 MHz) frequency, which is much easier to measure. Once the phase of the beat signal is detected, it can be used to digitally shift the phase of the recovered RF signal at the receiving end of the link. It is important to note that what is really measured in this method is the phase delay, not the group delay of the modulation which is ultimately what we are interested in. In order to compensate for group delay, a first-order correction – based on actual measurements of a given fiber type for a range of temperatures – is applied to the raw phase delay measurements.

0.1.3.2 Pulsed systems

In the pulsed system depicted in figure 6, a mode-locked laser is synchronized with the external RF signal, which determines its repetition rate. Narrow light pulses come out of the laser and travel through dispersion-compensated fiber to the receiver, where a fraction of the light is bounced back towards the emitter. An optical cross-correlator measures the degree of coincidence of the two pulse trains and the result is used to control a piezo actuator that changes the fiber length so as to keep a constant group delay. This mechanical actuator is unavoidable because the cross-correlator needs the pulses to overlap at least partially in order to give a meaningful reading. On the other hand, the pulsed system controls group delay directly, so no ad-hoc conversion between phase delay and group delay is needed.

0.1.4 Technical issues

Both types of systems have been successfully deployed over distances of several hundred meters, so the 25 km needed by CLIC is unknown territory. In the case of the CW system, one potential source of concern is the need for a model of phase-group delay corrections vs. temperature. If the temperature is not uniform over the complete length of the fiber, as can easily be the case in CLIC, a solution will need to be found.

Another potential issue is Brillouin scattering, a non-linear effect which is especially strong in optical fibers, due to the large optical intensities in the fiber core. Brillouin scattering in fibers leads to a limit of the optical power which can be transmitted, since above a certain power threshold, most of the light is scattered or reflected. This is especially a problem for narrow-band optical signals, and for long fiber lengths. In the pulsed system, Brillouin scattering is not much of an issue, since the power spectral density, which is the important quantity, is much smaller than for narrowband CW systems.

Synchronization of lasers with the optical reference is another field which needs further work. Currently, only synchronizing with an RF signal has been tested, but an all-optical system would be much more precise and avoid the extra optical-electrical conversion. On the optical source side, it is important to find lasers with a coherence length which will allow efficient stabilization on a 25-km fiber link.

For the pulsed system, a major difficulty will be the dispersion control of the optical fibers, as the fiber length increases. Due to the large dispersion in optical fibers, additional problems due to nonlinearities are not expected compared to shorter fiber links.

Another major difficulty in both schemes will be Polarization Mode Dispersion (PMD), which can

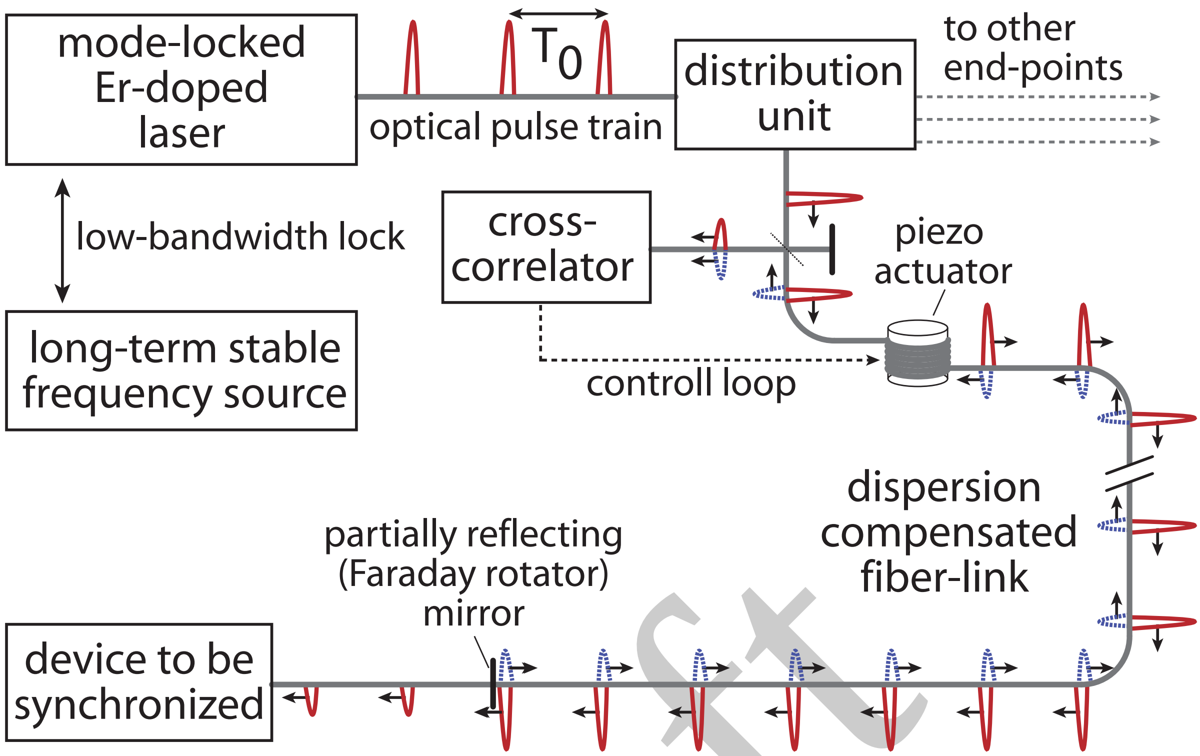

Fig. 6: Pulsed system for femtosecond timing distribution (courtesy of F. Loehl).

deteriorate the long-term timing stability especially for long link-lengths. The polarization state of light transmitted through a single mode fiber changes constantly, due to birefringence in the fiber. The larger this birefringence, the larger the propagation delay difference between the two polarization modes. Since the pulses or the CW light travel in different polarization states on both ways through the fiber, the travel times for both ways are not equal. The fiber birefringence and therefore the polarization state of the light further changes with temperature, which has the consequence that even if the light propagation delay for one round-trip is stabilized, the timing at the end of the link still shifts, since the propagation delays for both directions change in different ways.

Finally, one very important difference between the solutions currently deployed and the CLIC scenario is the scale of the project. Current implementations do not scale very well beyond some tens of destinations. This is mainly due to the number of ports typically available in optical components, such as splitters/combiners. For CLIC, one of the challenges will be to explore in detail the real needs of each destination and come up with a strategy for partitioning the system in such a way that a reasonable compromise between performance and cost can be found.

0.1.5 Cost considerations

The typical cost of a stabilized timing link is currently around 25 kEuro. Taking into account the hundreds of destinations needed in CLIC, it is very important to carefully study ways in which this cost can be driven down. CW systems have less opto-mechanical components and therefore have in principle the potential for being more cost-effective, but this remains to be studied in detail. Irrespectively of the technology used, one important aspect which can have a major impact on cost is whether the CLIC synchronization system can be partitioned in several loosely-coupled subsystems.

0.1.6 Outlook for Technical Design Report phase

As can be seen from the previous discussion, there are many open questions which remain to be answered before the TDR is published:

- The most important challenge is to make the existing solutions work with fiber lengths of around 25 km. This will involve work on the coherence length of laser sources and also an investigation of the problems induced by Brillouin scattering at high optical powers. A demonstrator of each technology needs to be built before an optimal decision can be taken for the best compromise between price and performance.
- Synchronizing remote lasers with the general timing system also needs to be investigated further, especially if an all-optical system is to be considered. In relation with this investigation, arrival-time monitors need to be studied since they are the most important diagnostics tool required to achieve good timing stability.
- A detailed study on the partitioning of the timing system among its different clients would potentially allow several semi-independent systems and have an important impact on costs. This study should be combined with an investigation of ways to extend the number of links supported by current solutions.

Bibliography

[1] Enrico Rubiola, Phase Noise and Frequency Stability in Oscillators, Cambridge University Press, 2009.

[2] Alan V. Oppenheim and Alan S. Willsky and S. Hamid Nawab, Signals and Systems, Prentice Hall, 2nd edition, 1996.

[3] Athanasios Papoulis and S. Unnikrishna Pillai, Probability, Random Variables and Stochastic Processes, McGraw Hill, 4th edition, 2002.

[4] Floyd M. Gardner, Phaselock Techniques, Wiley, 3rd edition, 2005.

[5] Erk Jensen, Drive Beam RF System, CLIC workshop 2009.

[6] Thibaut Lefevre, Progress on Instrumentation, International Workshop on Linear Colliders 2010.

[7] Daniel Schulte, Concept of MB-DB phase alignment and status of solutions, International Workshop on Linear Colliders 2010.

[8] Florian Loehl, Linac timing, synchronization and active stabilization, PAC 2011.