Possible Staging Scenarios for the ILC

Benno List[1]

[1]DESY, 22603 Hamburg, Germany

The ILC baseline design describes an accelerator with 500 GeV centre-of-mass energy. In view of the recent discovery of a Higgs particle, possible scenarios for a staging of the centre-of-mass energy and luminosity are presented, including the possibility to run the electron and positron beams at different energies in order avoid the need of a dedicated electron pulse for positron production at centre-of-mass energies below 250 GeV. Dedicated studies to determine the most efficient beam energies for the ILC physics program will help to optimize the machine layout, and are thus needed now, before a decision on the energy reach and possible intermediate stages will be met. A figure of merit is proposed that quantifies the value of an amount of integrated luminosity in terms of physics performance.

1 Introduction

The recent discovery [1, 2] and first characterization [3, 4] of "a" Higgs boson with a mass of 126 GeV at LHC has worldwide renewed the interest to build the ILC [5]. Measurements of the Higgs properties, such as couplings, width, and quantum numbers, can be performed under the ILC's clean experimental conditions with very good accuracy [6]; some crucial measurements, such as the absolute coupling to the Z^0 that sets the scale for all partial decay widths, are only possible at the ILC.

Precision measurements of the Higgs and top quark properties provide a rich physics program for the ILC [7], spanning the energies between 230-270 and 500-600 GeV. Recently, discussions have started about a scenario in which the ILC would be built in two or more stages, with rising beam energies, in the hope to reduce the initial investment, and possibly the construction time.

The ILC baseline design is a blueprint for an accelerator with a 500 GeV centre-of-mass energy, which could easily be adapted to yield a machine with 550 or 600 GeV CME. The baseline design foresees a possible later extension to 1 TeV centre-of-mass energy, and the central part, in particular the beam delivery system (BDS) is designed to accommodate the increased beam energy. However, no detailed design for this extension is yet available, and it is assumed that an energy upgrade would be designed at a time when further progress in superconducting cavity technology will make it possible to produce and operate cavities at larger gradients than currently possible. The TeV upgrade of the ILC is beyond the scope of this report; more information can be found in the Technical Design Report (TDR) [8].

In the following, some issues from the accelerator point of view will be presented first that may shed some light on the possible advantages and disadvantages of the scenarios under discussion. Then the problems of running at low energies (at 230-270 GeV, slightly above the threshold for Z^0h production) are discussed, with asymmetric running as a possible solution. Section 3 presents prospects to increase the luminosity, then some staging scenarios that have been investigated in somewhat more detail are presented. Section 5 is concerned with the need to define a figure of merit that would allow a more quantitative comparison of different running options, and thus help to optimize the machine design.

2 Accelerator Issues

2.1 Positron Source

The production of positrons in sufficient quantities is a key issue for the ILC operation. The baseline solution is a source where gamma ray photons are produced off the main electron beam in a superconducting helical undulator section. These photons impinge on a rotating target, where positrons are produced and subsequently collected by a flux concentrator. Photon production by a helical rather than a more common planar undulator offers two advantages: the photon intensity is twice as large, and the photons are longitudinally polarized, which results in a longitudinal polarization of the positrons of up to 60 %. All three parts, the undulator, target, and flux concentrator, pose significant engineering challenges. Therefore the baseline design includes a 50 % overproduction margin, i.e. if everything works according to plan, 50 % more positrons can be delivered to the damping rings than needed.

The energy spectrum and overall intensity of the photons produced in the helical undulator, and thus the resulting positron yield at the target, depends strongly on the beam energy. The baseline design foresees 42 undulators of 3.5 m length, for a total active length of 147 m. However, 24 additional undulators could be installed, bringing the total length to 231 m. In the baseline configuration the photon intensity is sufficient for the nominal yield[1] of 1.5 for electron beam energies above 150 GeV. Below that energy, the positron yield drops rapidly, such that at 125 GeV, the positron intensity is half that at 150 GeV. Part of that loss could be compensated, if necessary, by the installation of more undulator modules, but the fact remains that below 125 GeV electron beam energy positron production with the helical undulator source becomes rapidly unpractical.

The solution to this problem is the so-called "10 Hz" scheme (see Tab. 1), in which alternating electron pulses of 150 GeV for positron production and at lower energies for collisions are produced. This scheme essentially allows to provide electron beams for physics down to energies as low as 45 GeV needed for Z^0 running[2], albeit at the price of producing (and thereafter wasting) an additional 150 GeV beam. More implications of the 10 Hz scheme are discussed below.

2.2 Damping Rings

One main challenge for the damping rings is to reduce the vertical emittance of the beams, in particular the positron beam, in a very short time: the initial emittance of the positrons is approximately $\epsilon_\mathrm{y} \approx 0.8\,\mu$m, while the final vertical emittance is 10 pm rad, almost six orders of magnitude smaller. This reduction has to be achieved within $t = 200$ or even 100 ms, which means that the vertical damping time τ_y has to of the order of $\tau_\mathrm{y} \leq 2t/\ln(\epsilon_\mathrm{y,ini}/\epsilon_\mathrm{y,final}) = 15.5$ ms for 10 Hz operation [11]. The large energy loss per turn that is necessary for the required synchrotron radiation damping is achieved by the insertion of 54 wigglers in each ring, which lead to an energy loss of up to 7.7 MV per turn and a vertical damping time of $\tau_\mathrm{y} = 13$ ms. This results in up to 3.3 MW RF power necessary to store the a positron beam at the design current of 390 mA, which can be compared to the average beam power of the accelerated positron beam, which is 2.6 MW for a 250 GeV positron beam.

When the number of bunches per pulse was halved in the SB2009 process in order to save costs, the circumference of the damping rings was also halved, so that the bunch spacing and average current did not change. Doubling the number of bunches again would increase the current in each damping ring to 780 mA, which is feasible for the electron ring, but difficult for the positron rings, which would suffer from electron cloud formation and resulting beam instabilities. Therefore, the damping ring tunnel design leaves space for the installation of a second positron damping ring above the electron ring (the first positron damping ring lies below the electron ring), so that each ring would carry the same number of bunches as in the baseline configuration.

2.3 Main Linac

The performance of the Main Linac is limited in several places, some of which we will discuss in the following.

The most notorious performance limit is the accelerating gradient g, which determines how much energy per unit length can be transferred to the beam[3]. The ILC baseline assumes a gradient, averaged over all cavities, of $g = 31.5$ MV/m, with a spread of the gradients of individual cavities within ±20 %. The gradient of a cavity is measured during fabrication in a vertical test stand. There, the cavities have to achieve a gradient of 35 MV/m on average; the 10 % reduction accounts for the expected performance loss between the performance in the vertical test stand and within a cryomodule that houses eight or nine cavities.

[1] The positron yield is defined as the ratio of positron bunch intensity at the damping ring entrance divided by the electron bunch intensity at the undulator.

[2] Note that 45 GeV running for physics production is not included in the official baseline design.

[3] The maximal electric field in the cavities is larger by a factor $\pi/2$.

Increasing the accelerating gradient has been the goal—and the result—of more than twenty years of successful R&D [13]. Increasing the gradient means that the same beam energy can be achieved with fewer cavities, fewer cryomodules, and a shorter tunnel, in short, at reduced cost.

However, not everything gets cheaper with higher gradient; two quantities scale quadratically with gradient, and thus rise linearly with gradient for a linac with a given total beam energy: the stored energy in the cavity, and its dynamic heat load.

It is well known that the energy density of any electric field, and thus the energy stored within a resonating cavity, grows quadratically with the field strength. Because the ILC operates in a pulsed mode, at each pulse this energy has to be transferred into the cavity, and after the pulse it has to be dumped to a load[4] at room temperature.

At the ILC in the baseline configuration, a full RF pulse lasts 1.65 ms, of which 0.93 ms (56 %) are spent to fill the cavities, and only 0.73 ms (44 %) to accelerate the beam. The corresponding numbers for a luminosity upgrade with twice the number of bunches is 0.61 ms (39 %) for filling and 0.96 ms (61 %) for acceleration.

Thus, from an efficiency point of view it is desirable to accelerate more bunches per pulse, in order to make optimal use of the stored energy in the cavity, which is lost at the end of each pulse. The limits to that are posed by the maximum pulse length of the klystrons of about 1.6 ms, and the number of bunches that can be stored in the damping rings. In addition, the total electrical energy that is needed for one beam pulse has to be stored in the modulators, so a larger pulse energy means more modulators.

The second quantity that grows quadratically with gradient is the dynamical heat load. Although superconducting materials transport DC current losslessly, alternating fields penetrate the superconducting material within a finite skin depth, and there they accelerate the unpaired electrons which generate resistance and thus heat. The power P dissipated within one cavity of length L_{cav} ($L_{\mathrm{cav}} = 1.038\,\mathrm{m}$ for the ILC) is given by [14]

$$P = \frac{g^2 L_{\mathrm{cav}}}{(r/Q) Q_0},$$

where $r/Q = 1036\,\Omega$ is the shunt impedance per unit length, which depends only on the shape of the cavity, and Q_0 is the cavity's quality, which depends on the surface resistance of the cavity material. Therefore, the dynamic heat load from cavity heating (which accounts for 76 % of the cryomodule heat load at 2 K and 45 % of the overall heat load [15]) grows linearly with the cavity gradient, if the overall beam energy (and hence the product of the number of cavities times the gradient) is kept constant.

In summary, higher gradients reduce costs for tunnels, cavity material, cavities, couplers, and cryomodules, but increase costs for RF and cryogenic equipment and power consumption. Hence, there must exist a gradient that optimizes the overall cost, which has been estimated for the ILC to be beyond 60 MV/m [16]

2.4 Beam Delivery System

The main tasks of the beam delivery system are the measurement of beam properties, the collimation of the beams, and the final focussing. Measurement of the beam properties include energy, polarization and emittance; these measurements are used to tune the machine and correct effects such as coupling of vertical and horizon betatron oscillations, but also to dump beams to protect the detector, e.g. after a klystron trip. The collimation section removes halo particles with large amplitudes of the transverse (betatron) oscillations or large deviations from the nominal energy. The energy collimation section consists of two consecutive bends with zero net angle (a so-called dogleg) that displace the beam laterally. In the middle of the dogleg, off-momentum particles are displaced laterally from the main beam (this energy-dependent displacement is called the dispersion) and can be collimated. However, strong bending fields in regions of large dispersions cause emittance growth and thus have to be avoided; this effect grows fast with energy. Thus the allowed emittance increase limits the bending angles of the magnets employed in the dogleg and determine its length. The ILC design foresees a 1100 m long dogleg that is long enough to accommodate a 500 GeV beam (for 1 TeV centre-of-mass energy) with acceptable emittance growth. In the baseline configuration with 250 GeV beams, only every fifth magnet will be installed, which is sufficient at this energy.

[4]When the klystron is switched off, and there is no beam in the cavity, the RF field leaves the cavity via the waveguide system and is absorbed in a special, water-cooled load.

A basic property of longitudinal acceleration of beams is that it leaves the normalized emittance $\gamma\epsilon$ constant[5], where γ is the Lorentz boost of the accelerated particle, so that the emittance ϵ decreases with energy E as $1/\gamma$, or $1/E$. Consequently, for a given value β^* at the interaction point, the RMS beam size $\sigma = \sqrt{\beta^* \epsilon}$ also decreases, and the luminosity, which is proportional to $1/(\sigma_x \sigma_y)$, rises proportional to E. However, β^* is not only limited by the achievable strength of the final focus quadrupoles, which makes focussing harder at large beam energies, but also by the beam size within these quadrupoles, which has to fit into the magnet's aperture. This beam size limits the vertical beam size at lower energies, which means that the luminosity drops somewhat faster the with E if the beam energy is reduced, as can be seen in Tab. 1.

3 Running at Low Energies

3.1 The 10 Hz Scheme

As discussed above, the 10 Hz scheme is based on decoupling the two functions of the electron beam, namely to collide with the positrons for physics measurements, and to produce photons for the positron source. Originally this scheme was proposed for a machine capable of running at 250 GeV beam energy or more, to allow such a machine to operate below 125 GeV beam energy. Under that circumstances, the Main Linac would run at approximately half its nominal gradient, and therefore need about half of the maximum power for RF production, and even less for cryogenic cooling: the cryogenic power needed for the Main Linac is dominated by dynamic losses, i.e. by the heat deposited in the superconducting cavities by the accelerating fields, an those grow quadratically with the gradient. Therefore, an electron Main Linac capable of delivering electron pulses of 250 GeV at a repetition rate of 5 Hz is able to deliver alternating pulses of 150 GeV and $\leq$ 125 GeV at 5 + 5 Hz without the same cryogenic cooling power and the same overall electrical power consumption. However, although no additional investments (such as larger cryo plants or larger transformer stations) in the Main Linac are needed to enable the operation under the 10 Hz scheme, the total efficiency (electrical energy, or dollars, per unit of integrated luminosity) is significantly reduced by the necessity to accelerate a second high energy electron beam solely for positron production. In the case of Z^0 running, the electron beam energy used for positron production would be a factor 1.6 larger than the total beam energy used for collisions.

For a staged machine, the situation becomes more complex, and even less attractive: While a symmetric machine (where electrons and positron beams have the same energies) operating at 250 GeV centre-of-mass energy needs a 125 GeV electron beam, efficient positron production requires a machine capable of delivering a 150 GeV electron beam. Even worse, now the electron linac has to run at full gradient for the positron production beam and the "physics" beam, so that it requires almost twice the electrical power and cryogenic cooling capacity to allow running with the 10 Hz scheme, compared to nominal 5 Hz operation. While this may make running with alternating beams for the purpose of calibration runs at the Z^0 possible, it appears highly undesirable for real physics operation.

For the damping rings, the 10 Hz operation also poses significant challenges, because the lingering time of the beams in the damping rings is halved from 200 ms to 100 ms. Therefore, the already challenging task to reduce the positron vertical emittance by more than five orders of magnitude within 200 ms now becomes twice as hard. To achieve this goal, the vertical damping time has to be reduced from 24 ms to 13 ms, which is achieved by adding more wigglers (for more synchrotron radiation damping) and more RF power to replenish the synchrotron radiation loss. It should be noted that this means that the damping rings in the new configuration are capable of delivering fully damped beams at twice the rate that is needed for nominal physics operation, which opens up the possibility to increase the luminosity of the machine by an increased repetition rate.

In summary, any staging scenario should avoid physics data taking for any extended time in the 10 Hz scheme at full gradient, because the cost to luminosity ratio becomes unfavorable. Running at sufficiently large electron beam energies (above 125 GeV) and possibly an extended helical undulator will help to avoid 10 Hz running. However, running at higher pulse rates to increase overall luminosity may be possible.

[5]A slight rise of $\gamma\epsilon$ along the main linac is caused by imperfections and the need to follow the earth's curvature.

3.2 Asymmetric Running

An alternative way to run at lower centre-of-mass energies is the operation with beams of different energies. For instance, a centre-of-mass energy $E_{\mathrm{CM}} = 235\,\mathrm{GeV}$ could be achieved by colliding 150 GeV electrons with 92 GeV positrons, with a moderate boost of $\beta = 0.25$.

Detailed studies about the ramifications of such a scheme are still outstanding. Obvious questions to be answered from the detector and analysis point of view are how the resulting boost would affect the acceptance and resolution of tracks and jets. Since the Higgs strahlung process $e^+e^- \to Z^0h$ that provides the motivation to run at energies around $230-270\,\mathrm{GeV}$ is asymmetric in itself, a moderate longitudinal boost should not be too problematic, in contrast to measurements at the Z^0 pole that involve forward-backward asymmetries. On the accelerator physics side, studies would be needed to figure out which luminosity would be the achievable in such a scenario, and what the resulting beam disruption parameters and backgrounds (e.g. from pair production) would be.

4 Increasing the Luminosity

An important change in the baseline parameters that occurred between the ILC Reference Design Report (RDR) of 2007 and the Technical Design Report (TDR) of 2013 was the reduction of the number of bunches per pulse from $N_{\mathrm{bunch}} = 2625$ to 1312 in oder to reduce the necessary RF power and thus save costs. During the beam pulse, the total RF power that has to be provided by the klystrons is simply given by $P = E_{\mathrm{beam}}/eI_{\mathrm{beam}}$, where E_{beam} is the total beam energy, $I_{\mathrm{pulse}} = Q_{\mathrm{b}}/\Delta t_{\mathrm{b}}$ the beam current during the pulse (with the bunch charge $Q_{\mathrm{b}} = en_{\mathrm{b}}$ and time between bunches Δt_{b}), and e is the elementary charge (see Tab.1 for the actual values).

Naively one would expect that halving the number of bunches but keeping the total pulse length $t_{\mathrm{pulse}} = n_{\mathrm{b}}\Delta t_{\mathrm{b}}$ constant, which halves also the bunch current, would allow a reduction of the number of klystrons by a factor of two as well. However, this neglects the fact the fill time t_{fill}, which is the time needed to ramp up the cavity gradient from zero to the nominal 31.5 MV/m, also increases, so that the maximum pulse time of 1.6 ms that can be provided by the klystrons is exceeded. Thus, while the number of bunches was halved, the current was reduced only by 36 % from 9 to 5.8 mA, and the number of the klystrons was reduced by one third.

Turning that calculation around shows that to double the luminosity by a doubling of the number of bunches per pulse requires only 50 % more klystrons, and improves the luminosity-to-power ratio significantly[6]. The damping rings are designed to accommodate this increased number of bunches. For the electron damping ring, it is expected that it can store the full number of bunches. In the case of the positron damping ring, a doubling of the number of bunches in the ring will necessitate to build a second positron damping ring, which is foreseen in the tunnel design.

A further increase of the number of bunches per pulse is limited by capacity of the damping rings due to the onset of instabilities, and by the maximal pulse length that the klystrons can provide. However, the damping rings are designed to achieve the necessary damping within 100 ms, which would allow a doubling of the luminosity by doubling the pulse rate to 10 Hz, provided that the cryogenic plants are upgraded to provide more cooling capacity. Whether the positron source, in particular the target, is capable to run at 10 Hz needs, however, to be checked. What makes a doubling of the pulse frequency attractive is the fact that no additional klystrons and modulators are needed; the modulators (which accumulate the energy needed to provide the high power pulses for the klystrons) would simply charge at twice the rate. Such a luminosity upgrade could be realized relatively fast, provided that the cryogenic plants are can designed for easy upgradeability. The downside of this upgrade path is that the luminosity-to-cost ratio hardly improves.

Another advantage of a luminosity increase through an increased repetition rate is that it can be continuously adjusted: any repetition rate up to the maximal rate of 10 Hz is feasible, whereas the number of klystrons can only be increased in fixed steps. In a staged scenario it is conceivable that one would install the full cryogenic capacity early on, and run the accelerator at each energy at the maximum repetition rate allowed by the available cooling capacity.

[6] Even more so as many other contributors to the power consumption are independent of the beam current.

			baseline				upg.
Centre-of-mass energy	E_{CM}	GeV	200	250	350	500	500
Luminosity pulse repetition rate		Hz	5	5	5	5	5
Positron production mode			10 Hz	10 Hz	nom.	nom.	nom.
Site AC power consumption	P_{AC}	MW	114	122	121	163	206
Average beam power	P_{ave}	MW	4.2	5.3	7.4	10.5	21.0
Bunch population	N	$\times 10^{10}$	2	2	2	2	2
Number of bunches	n_b		1312	1312	1312	1312	2625
Linac bunch interval	Δt_b	ns	554	554	554	554	366
Beam current in pulse	I_{pulse}	mA	5.8	5.8	5.8	5.8	8.8
Beam pulse duration	t_{pulse}	µs	727	727	727	727	961
Average gradient	g	MV/m	12.6	15.8	22.1	31.5	31.5
RF pulse length	t_{RF}	ms	1.10	1.19	1.37	1.65	1.57
RMS bunch length	σ_z	µm	300	300	300	300	300
Normalized horizontal emittance at IP	$\gamma\epsilon_x$	µm	10	10	10	10	10
Normalized vertical emittance at IP	$\gamma\epsilon_y$	nm	35	35	35	35	35
Horizontal beta function at IP	β_x^*	mm	16	13	16	11	11
Vertical beta function at IP	β_y^*	mm	0.34	0.41	0.34	0.48	0.48
RMS horizontal beam size at IP	σ_x^*	nm	904	729	684	474	474
RMS vertical beam size at IP	σ_y^*	nm	7.8	7.7	5.9	5.9	5.9
Vertical disruption parameter	D_y		24.3	24.5	24.3	24.6	24.6
Fract. RMS energy loss to beamstrahlung	δ_{BS}	%	0.65	0.97	1.9	4.5	4.5
Luminosity	L	10^{34} cm^{-2} s^{-1}	0.56	0.75	1.0	1.8	3.6
Fraction of L in top 1% E_{CM}	$L_{0.01}$	%	91	87	77	58	58
Electron polarisation	P_-	%	80	80	80	80	80
Positron polarisation	P_+	%	30	30	30	30	30
Electron relative energy spread at IP	$\Delta p/p$	%	0.20	0.19	0.16	0.13	0.13
Positron relative energy spread at IP	$\Delta p/p$	%	0.19	0.15	0.10	0.07	0.07
Beamstrahlung parameter (av.)	Υ_{ave}		0.013	0.020	0.030	0.062	0.062
Beamstrahlung parameter (max.)	Υ_{max}		0.031	0.048	0.072	0.146	0.146
Energy loss from BS	δE_{BS}	%	0.65	0.97	1.9	4.5	4.5
e^+e^- pairs per bunch crossing	n_{pairs}	10^3	45	62	94	139	139
Pair energy per B.C.	E_{pairs}	TeV	25	47	115	344	344

Table 1: Summary table of the 200–500 GeV baseline parameters for the ILC [8, 9, 10], including parameters for a possible luminosity upgrade. The numbers for energies between 200 and 350 GeV correspond to a situation where the full 500 GeV accelerator is being operated at a reduced energy.

			Minimal	Scenario 1	Scenario 2	Baseline
Centre-of-mass energy	E_{CM}	GeV	250	250	250	500
... extensible in tunnel to	E_{CM}	GeV	250	500	500	500[a]
Site AC power consumption	P_{AC}	MW	120[b]	120[b]	125[b]	163
Relative cost (estimated)		%	67[b]	73[b]	75[b]	100

Table 2: Overview over possible staging scenarios [17].

[a] The baseline design includes an option to extend the energy to 1 TeV, but not within the initial tunnel.

[b] If the 10 Hz scheme is avoided, an estimated 25 MW of AC power and 3 % of costs can be saved.

5 Staging Scenarios being Discussed

Fig. 1 illustrates several possible scenarios for an ILC with an (initial) energy of 250 GeV, in comparison with the nominal 500 GeV baseline design [17].

A minimal Higgs factory, limited to about 250 GeV centre-of-mass energy, is one possible option. In such a scenario, the central region could be reduced, because the BDS would not have to accommodate a high-energy beam, which would save some cost. Altogether, such a machine is estimated [17] to cost approximately 67 % of the machine in the full TDR baseline design configuration.

A truly staged machine, with a central region that makes future upgrades possible, could come in two forms:

In scenario 1, the initial stage would include a main tunnel long enough to accommodate the full accelerator, but have the bunch compressor immediately adjacent to the shortened main linac, so that an energy upgrade necessitates a relocation of the RTML turn-around and bunch compressor, which corresponds to about 2200 m of beamline on either side. This scenario is estimated to cost about 73 % of the full machine. The moderate price difference compared to the minimal Higgs factory illustrates that the tunnel itself is not the biggest cost driver at the ILC. A possible advantage of this scenario might be that not the full tunnel has to be ready from the beginning on; however, in the Japanese site, the tunnel sections would be excavated by several teams working in parallel, so that a shorter tunnel would not be finished much faster than the full tunnel.

In scenario 2, the turn-around and the bunch compressor are built at their final positions from the beginning on, and the space reserved for the installation of further cryomodules is bridged by a transfer line in the first stage. This option is estimated to cost about 75 % of the full machine. This scenario would make upgrading much faster and cheaper (no relocation of the turnaround, no re-commissioning of the bunch compressor), and conceivably allow to upgrade the energy with an intermediate step around the top threshold at 350 GeV. If such a scenario is envisioned, the cryomodule production schedule could be stretched, which may actually save costs because less infrastructure for cavity and coupler production, cryomodule assembly, and testing would be needed.

6 Optimizing the Energy: An Experimental Figure of Merit

Defining the (initial) energy reach of the ILC will depend heavily on factors outside the control of the physics community, such as constraints from accelerator physics, constraints from the chosen site, available funding, and politics. One constraints from accelerator physics comes from the fact that the round trip that a positron takes from (roughly) the middle of the Damping Ring injection line to the IP must be an integer multiple of the DR circumference of 3.2 km, which results in a quantization of the (positron) beam energy in steps of roughly 34 GeV, assuming that the available tunnel space will be fully equipped with cryomodules. The site may have an influence because some places may be more or less suitable for access shafts. Financial considerations always play a role, obviously, because a larger linac costs more, but also because initial investments can be partially traded against running costs: a longer machine costs more initially, but may deliver more events per unit running cost.

However, the physics and detector community can certainly make an important input to these considerations, if a quantitative figure of merit can be given that indicates how much luminosity at a given centre-of-mass energy (and other conditions, such as beamstrahlung or disruption parameter) is needed to obtain a certain physical result.

For example: Consider the measurement of the Higgs production cross section with the recoil mass technique in $e^+e^- \to Z^0 h \to \mu^+\mu^- X$, where the Higgs mass peak is reconstructed from the $\mu^+\mu^-$ recoil mass spectrum. We can expect the relative statistical uncertainty $\delta\sigma/\sigma$ of the cross section σ to be given by $\delta\sigma/\sigma = q/\sqrt{\sigma L}$, where L is the integrated luminosity, and q is a quality factor that depends on the performance of the detector and analysis code, but also on the centre-of-mass energy and the beam parameters, in this particular case on the beam energy spread (or, equivalently, the disruption parameter).

Thus, the integrated luminosity L that is needed to achieve a certain statistical precision $\delta\sigma$ is given by $L = q^2/\sigma/(\delta\sigma/\sigma)^2$, so that q^2/σ serves as a figure of merit for different beam energies and running conditions.

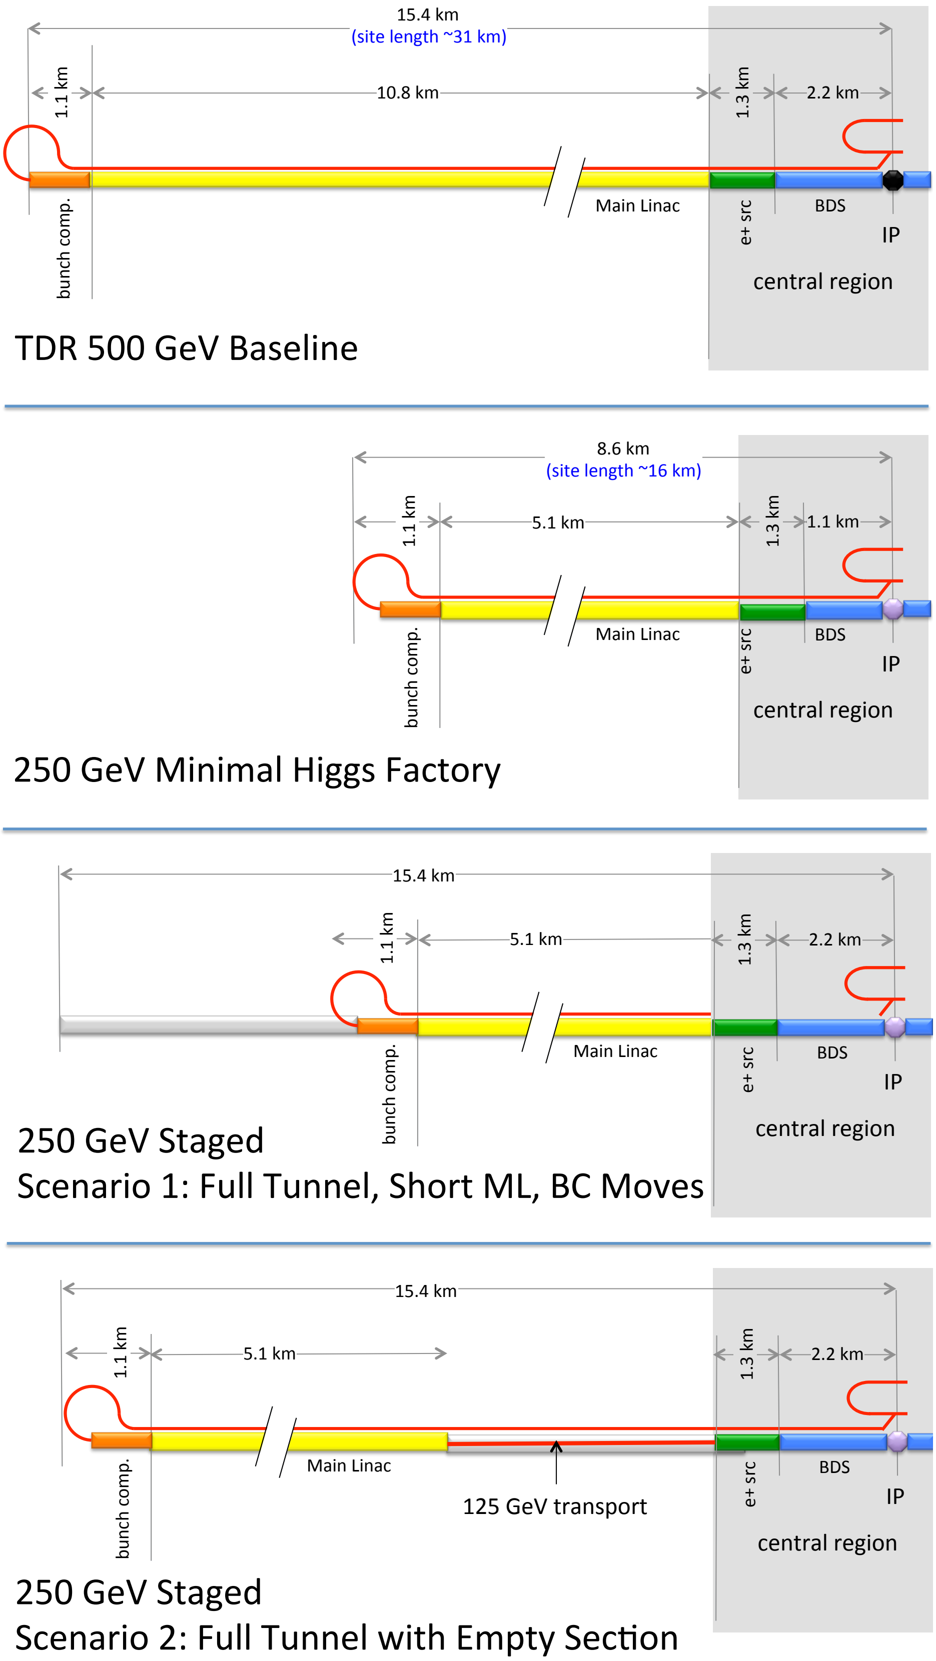

Figure 1: Staging scenarios [17].

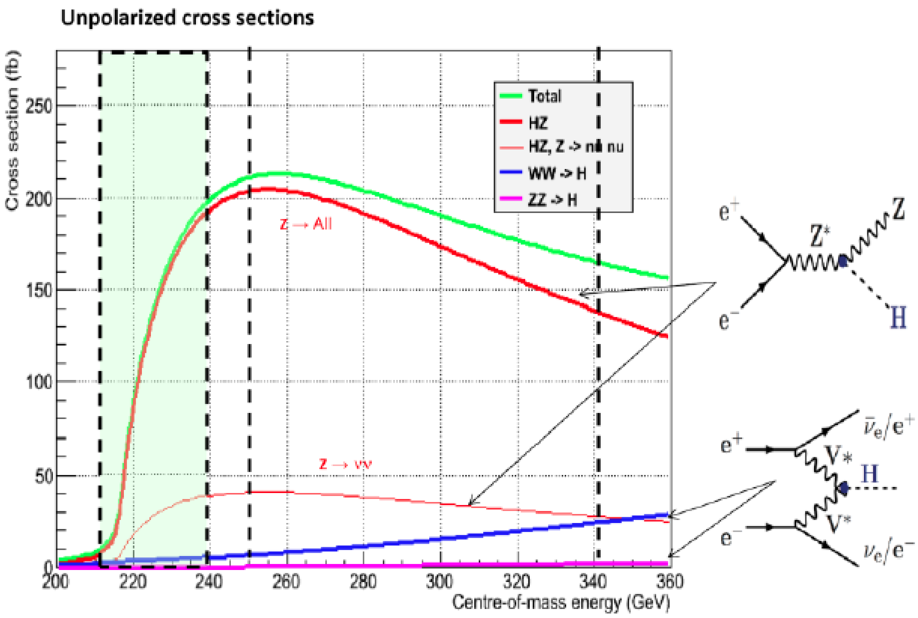

Figure 2: Higgs boson production cross section [19].

For the recoil mass measurement, the mass resolution is best around 20 GeV above the threshold for Z^0h production [18], i.e. around 235 GeV (see Fig. 2). However, below 250 GeV or so, the helical undulatory becomes so inefficient that it may be necessary to resort to the 10 Hz running scheme, where between two beam pulses with colliding e^+e^- beams (that occur at 5 Hz) an intermediate e^- beam pulse is accelerated to 150 GeV just for positron production. Thus, one 150 GeV beam alternates with two beams of 235 GeV total beam energy, therefore only $235/(150 + 235) = 61\,\%$ of the beam energy are used for physics, with a corresponding (though not necessarily proportional) increase in costs per inverse femtobarn of produced integrated luminosity. In addition, the instantaneous luminosity (at constant beam disruption) of a linear collider rises at least linearly with the beam energy, because the beams get smaller at larger energies. In essence, the operation cost per inverse femtobarn falls with energy. Therefore, from an accelerator physics point of view the conditions to produce luminosity become more favorable as the centre-of-mass energy is increased to, say, 260 or 270 GeV. This effect is further enhanced by the fact that the Z^0h cross section has its maximum around 260 GeV.

Experimentally, on the other hand, conditions become less favorable with increasing energy because the boost of the Z^0 in the laboratory frame rises, which leads to a deterioration of the recoil mass resolution, and thus to more background under the broadening peak and a reduced quality factor q. However, existing studies indicate [18] that the measurement of the Z^0h coupling, which is probably the more relevant measurement to be performed with the recoil mass method, works well at 250 GeV and is much less affected by an increase of the centre-of-mass energy than the mass measurement.

A quantitative investigation of the pros and cons of running at higher or lower centre-of-mass energies is of interest now, before the layout of the detectors and the accelerator has been finalized, because it may have an impact on this design: If it turned out that indeed energies significantly below 250 GeV are necessary to achieve the physics goals, then an optimization of the helical undulatory positron source should be considered, for instance with a longer undulatory or an undulatory with different parameters. If, on the other hand, running at higher energies such as 250 or 270 GeV turned out to be viable, then further parameter changes might result: The reduced resolution in the recoil mass spectrum would result in a smaller impact of the beam energy spread, allowing running at larger disruption parameters, which increases the luminosity. Since the recoil mass resolution is the driving force for the tracking resolution of the detectors, even that might be influenced.

Similar considerations should be applied in the determination of the maximum beam energy of the

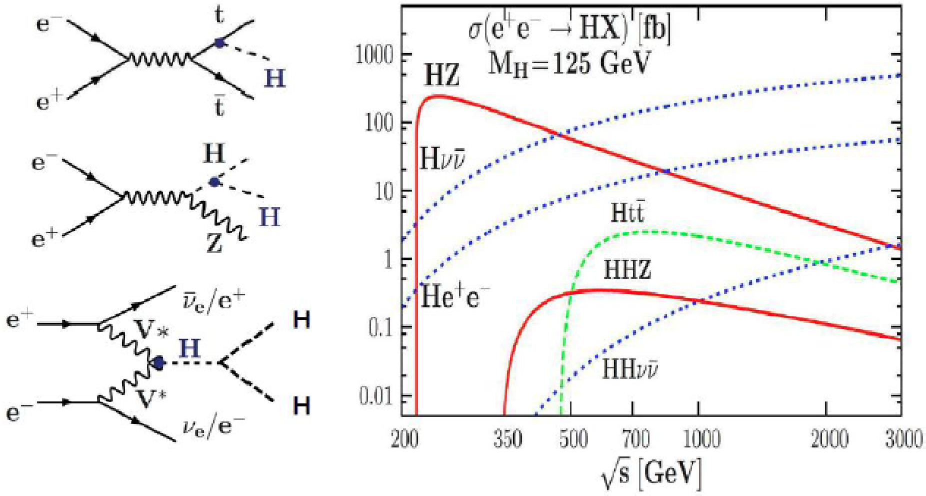

Figure 3: Higgs boson production cross section at larger energies [19].

accelerator (not of the possible TeV upgrade, but of the half-TeV-ish machine): From the point of view of Higgs physics, two measurements stick out that require energies around or above half a TeV: The top-Higgs-coupling, and the Higgs self coupling. The top-Higgs-coupling is based on the process $e^+e^- \to t\bar{t}h$ with a threshold of 475 GeV, with a cross section that rises rapidly between threshold and 600 GeV, as shown in Fig. 3. For this process, more energy is clearly better, so the trade-off is clearly between initial costs and operating costs. But again, to quantify the cost advantage that a rise of the centre-of-mass energy from 500 to, say, 550 or 600 GeV would have, experimental studies of the corresponding quality factor q (or $q^2\sigma$) are needed. In fact it may turn out that a measurement of the $t\bar{t}h$ coupling is unviable (meaning that it produces no noticeable improvement over the expected LHC performance) at 500 GeV, but viable at a slightly larger energy. For instance, the $t\bar{t}h$ cross section rises by a factor of 3.7 between 500 and 550 GeV [20]. From the machine point side, a quantitative evaluation of the relative size of running costs and investment costs might be interesting. One can guess that one year of machine operation costs would be comparable to the cost increase needed to raise the machine's energy by anywhere between 25 and 100 GeV.

Another process that profits a lot from increased centre-of-mass energy is the measurement of the Higgs self coupling. This measurement is based on the detection of double Higgs production events from a $h \to hh$ branching. As in single Higgs production, the intermediate Higgs can be produced by Higgs-strahlung off a Z^0, or by the fusion process $W^+W^- \to h$. However, other diagrams that do not involve a tri-Higgs coupling can also produce Higgs pairs, in particular repeated Higgs-strahlung off a Z^0 or a t-channel W. As Fig. 3 shows, the relevant cross sections are quite small, but in particular the $hh\nu\bar{\nu}$ cross section rises significantly with the centre-of-mass energy. Studies performed for the TDR [21] assume an integrated luminosity of the order of $2\,\text{ab}^{-1}$ at 550 GeV, corresponding to many years of running, and predict just a $5\,\sigma$ effect for that amount of data. If one estimates that one year of ILC operation might cost about as much as 50 GeV worth of Main Linac, it is well possible that at 600 GeV the prospects for this extremely important measurement might be better.

7 Conclusions

With the discovery of a Higgs boson at LHC, an exiting physics program for the ILC lies ahead of us. This program needs the ILC's capabilities at comparatively low energies around 230-270 GeV as well as at high energies around 500–600 GeV. Physics studies and deliberations among the physics communities which amounts of data at which energies and in which sequence would produce the most interesting and relevant physics results are needed now, in order to define the optimal final and intermediate configuration of the ILC. In several places one may trade initial investments for running costs to achieve an optimal physics performance at a given budget, within a reasonable time.

8 Acknowledgments

I thank Nick Walker, Jenny List, Eckhard Elsen, Karsten Büßer, Brian Foster, Georg Weiglein, Gudrid Moortgat-Pick and Sven Heinemeyer for many helpful discussions and explanations.

References

[1] G. Aad *et al.* [ATLAS Collaboration], Phys. Lett. B **716** (2012) 1-29 [arXiv:1207.7214 [hep-ex]].

[2] S. Chatrchyan *et al.* [CMS Collaboration], Phys. Lett. B **716** (2012) 30-61 [arXiv:1207.7235 [hep-ex]].

[3] The ATLAS Collaboration, *Combined coupling measurements of the Higgs-like boson with the ATLAS detector using up to* $25\,\mathrm{fb}^{-1}$ *of proton-proton collision data,* ATLAS-CONF-2013-034 (2013).

[4] S. Chatrchyan *et al.* [CMS Collaboration], *Observation of a new boson with mass near* $125\,GeV$ *in pp collisions at* $\sqrt{s} = 7$ *and* $8\,\mathrm{TeV}$, arXiv:1303.4571 [hep-ex] (2013).

[5] S. Heinemeyer *et al.* (eds), *Implications of LHC results for TeV-scale physics: signals of electroweak symmetry breaking,* Submitted to the Open Symposium of the European Strategy Preparatory Group, 2012, `https://indico.cern.ch/contributionDisplay.py?contribId=131&confId=175067`.

[6] M. E. Peskin, *Comparison of LHC and ILC capabilities for Higgs boson coupling measurements,* arXiv:1207.2516 [hep-ph] (2012).

[7] J. E. Brau *et al.*, *The physics case for an* e^+e^- *linear collider,* arXiv:1210.0202 [hep-ex] (2012).

[8] P. Burrows *et al.* (eds.), *International Linear Collider Technical Design Report, Volume 3: Accelerator, Part II: Baseline Design,* DESY 13-062, 2013.

[9] N. Walker, *ILC parameters,* 2012, `http://ilc-edmsdirect.desy.de/ilc-edmsdirect/item.jsp?edmsid=D00000000925325`.

[10] *ILC Technical Design Documentation,* `http://www.linearcollider.org/ILC/GDE/technical-design-documentation`.

[11] A. Wolski, *Damping ring design and physics issues,* lectures given at the Fourth International Accelerator School for Linear Colliders, Beijing, 2009, `http://ilcagenda.linearcollider.org/conferenceOtherViews.py?view=standard&confId=3475`.

[12] M. A. Palmer and S. Guiducci, *ILC Damping Ring parameters,* 2012, `http://ilc-edmsdirect.desy.de/ilc-edmsdirect/item.jsp?edmsid=D00000000960955`.

[13] P. Burrows *et al.* (eds.), *International Linear Collider Technical Design Report, Volume 3: Accelerator, Part I: R&D in the Technical Design Phase,* DESY 13-062, 2013.

[14] H. Padamsee, *Designing superconducting cavities for accelerators,* Proc. CERN Accelerator School 2002, Erice, Italy, CERN-2004-008 (2004) 233-252.

[15] T. Peterson, *ILC cryogenic heat loads and cryoplant parameters,* 2012, `http://ilc-edmsdirect.desy.de/ilc-edmsdirect/item.jsp?edmsid=D00000000994395`.

[16] C. Adolphsen, *Review of machine parameter / cost relationships,* presented at ILCWS'11, Granada, Spain, `http://ilcagenda.linearcollider.org/conferenceOtherViews.py?view=standard&confId=5134`.

[17] N. Walker, *ILC as a Higgs factory,* talk at Accelerators for a Higgs Factory: Linear vs. Circular (HF2012), Fermilab, 2012, `https://indico.fnal.gov/conferenceOtherViews.py?view=standard&confId=5775`.

[18] H. Li *et al.* [ILD Design Study Group Collaboration], *HZ recoil mass and cross section analysis in ILD,* arXiv:1202.1439 [hep-ex] (2012).

[19] A. Blondel *et al.*, *Report of the ICFA Beam Dynamics Workshop 'Accelerators for a Higgs Factory: Linear vs. Circular' (HF2012),* arXiv:1302.3318 [physics.acc-ph] (2013).

[20] C. Farrell and A. H. Hoang, Phys. Rev. D **72** (2005) 014007 [arXiv:hep-ph/0504220].

[21] J. Brau *et al.* (eds.), *International Linear Collider Technical Design Report, Volume 2: Physics,* DESY 13-062, 2013.