Responses from the GDE Executive Committee (GDE-EC) to reviews, comments and remarks on the ILC baseline configuration choices

Original comments are in italics. Responses from the GDE-EC are in regular text.

A. Review of the strawman BCD by B. Richter

1. Physics time-scale, expected LHC results etc. strongly favour an adiabatic energy upgrade path (i.e. Option 1: full length tunnel)

The arguments for and against Option 1 (tunnel length for 1 TeV) and Option 3 (tunnel length for 0.5 TeV) are more political than technical, and were discussed during the Frascati workshop (see below, part E), and will no doubt be discussed again. The essential point is that it is difficult to justify a 1 TeV machine at this time, and it is also difficult to argue that the 500 GeV machine, which can be justified, is cost-optimized if Option 1 is chosen.

From a physics point of view, having more upgrade potential is always better. At this point, there is no evidence that even 1 TeV will be sufficient. Neither Option eliminates the possibility of going above 1 TeV, provided that sufficient length on the site is available. If > 1 TeV is required, the 1 TeV tunnel can be extended further upstream after filling the downstream empty part, with the same extra cost of moving some components as in the case of Option 3.

2. Extending the tunnel backwards for the energy upgrade (current baseline) will require a post-damping ring booster (5-10 GeV) for emittance preservation in the long (~10km) transfer line.

Given the small (~0.1%) energy spread from the damping ring, it is expected to be relatively straightforward to construct a simple transfer line that can preserve the emittance. Clearly beam based alignment will be required, but the techniques developed for the other beamlines – particularly the main linac where the initial energy spread is an order of magnitude larger – will be more than adequate. Attention must be made to having adequately specified diagnostics (BPMs).

Accelerating the beam before bunch compression will result in a higher energy spread after compression. The best approach is to extract the beam, do the spin rotation, transport to the end, perform the turn-around and then compress. It would be slightly better to do the spin rotation after the turn-around but this is additional length for a small gain.

3. The rich, broad spectrum physics programme and the questionable technical feasibility of a push-pull solution strongly favour two detectors with two IRs.

The GDE-EC has requested that a complete study of a single IR configuration be completed during the next few months. As part of this study, the feasibility of a push-pull scenario for two detectors will be evaluated.

4. Tunnel alignment: agrees with non-laser straight solution.

While B. Richter clearly states that he agrees with the recommendation to abandon the laser-straight solution, he infers that a straight-line segmented solution be adopted, requiring discrete achromatic bending sections at specified intervals. The current BCD is to have a continuous curvature, unless cost-driven site-specific issues dictate otherwise; this was adopted because it was perceived to be the cheapest option, and there is currently no evidence of beam dynamics issues (with the caveat that there is still much work to be done).

5. Positron source: unable to comment at this time due to complexity of the arguments (needs to be better informed)

6. Number of tunnels: clearly agrees with the two-tunnel solution being mandatory. Disagrees with the 'White Paper' on suggestions that a single-tunnel solution may be practical with the Marx modulator

Nan Phinney clarified the Marx Modulator single-tunnel comment at the Frascati meeting, pointing out that it was not the intention of the White Paper to suggest a single-tunnel solution if the Marx Modulator was adopted as baseline, but that a single-tunnel solution with the modulator in the tunnel is not feasible with the current bouncer modulator.

7. Accelerator gradient: need to understand if mass production of 31.5 MV/m cavities is feasible (affordable); understanding the impact of vacuum breach.

- Mass production: the mass production of some 16,000 cavities with an average gradient of 37 MV/m (*not* 31.5MV/m!) is probably the biggest R&D challenge facing the ILC community. Much relies on the current R&D on understanding and reducing the performance spread, by specifying and accurately monitoring the EP parameters; this work must necessarily come from us (the ILC community) and cannot be expected from industry. Asking industry to produce a *small* number of prototypes is clearly desirable, but the required investment in the infrastructure (including the necessary quality control measures) may prove prohibitive for such a small number. This is an unfortunate situation, which can only be mitigated by close collaboration of industry with the infrastructure currently being set up at the various test facilities around the world. In the short term (next six months to a year), we are faced with the problem of coming up with a realistic cost for the cavities. The TESLA TDR costs were based on BCP+1400° C bake for an average (operational) performance of 23.4 MV/m. While many DESY experts think that the price of BCP+1400° C is probably about the same as EP+800° C, there still remains the question of a possible cost increase associated with the higher performance. It is difficult to imagine that not being the case. This question will need to be resolved through cost studies in laboratories and industries during the next year, carried out as part of the development of the Reference Design Report (RDR).

- Vacuum issues: the potential venting of the linac to air and the associated (and potentially catastrophic) consequences have been the topic of discussions for many years. Several points are covered in TESLA 2002-06. The most likely accident is a vent of beam vacuum to helium, which is not such a disaster (this would, for example, be the primary result of beam damage). 'Collision damage' to any part of the RF distribution system is protected by the double window in the coupler. Clearly, we need to have isolation valve interlock systems to mitigate (localise) the damage, but the exact spacing is a cost versus risk issue. These problems are also important for the XFEL, and tests like those proposed by B. Richter will be performed (although the schedule for these tests in not currently known). Clearly we must design the main linac to reduce the risk of 'external' damage from 'human intervention' as much as possible, but B. Richter's comments are well taken.

B. Review of the strawman BCD by K. Oide

1. Luminosity parameters: proposed 'flexibility' in parameters can be interpreted as designing for a luminosity of ~5×10^{34} cm^{-2}s^{-1}. Should design for 2×10^{34} cm^{-2}s^{-1} (nominal) as nobody will complain if we do not achieve the luminosity from the beginning. (Historical note on single-parameter set in HEP machines.)

The concept of the 'parameter plane' has been proposed to provide risk reduction in ensuring that the design luminosity will be reached, given that we do not know how the machine will ultimately perform. One parameter can be successfully traded off against another in the event of some unforeseen bottleneck. Thus, the design is aimed not at a machine with $\mathcal{L}$~5×10^{34} cm^{-2}s^{-1}, but at a machine that can provide $\mathcal{L}$ ~2×10^{34} cm^{-2}s^{-1} with some confidence.

This having been said, the GDE-EC agrees with K. Oide that we are effectively designing a machine with a factor of 2.5 safety margin. We definitely need to review the cost impact of this safety margin. In addition, it should be noted that this machine will have no competition. If two machines are competing, like PEP-II and KEK-B, a factor of 2 in the luminosity surely matters. If one is trying to improve an existing machine, a factor of 2 is also important. But in the design stage of a unique machine, reliability/availability issues are much more important than a factor of 2 in the design luminosity. There can be a chance of loosing a factor 10 or 100 if these issues are not properly treated.

2. Cost reduction by clever RF distribution system.

Through a combination of slightly detuning each cavity and adjusting its external *Q*, it is possible to drive each individual cavity close to its limit. The klystron has to have enough overhead to keep the vector sum over all cavities constant over the pulse. How well this system works depends on the spread of the performance in the RF units. The calculations have been made for the XFEL, and they believe they can regain about 67% of the voltage 'lost' due to running the unit at the weakest performer with respect to the average performance. (This can be pushed to 100% in principle if the fast piezo-tuners are individually programmable.) There are many parameters to balance, and there may also be an associated small loss in efficiency. If we extrapolate this to the current baseline specification for cavity production, we might argue that we can achieve an average gradient of ~33 MV/m, or close to a 5% increase, which is comparable to the 'overhead' that is currently stated in the BCD.

While 5% is not negligible, it's not a very significant amount either. To realize this gain, one must deal with the complexity of individually tuning each cavity, and ensuring the stability of the resulting balance. The tuning procedure must be automated over the entire linac (it is certainly not practical to do this by hand!) and this requires study.

Because of the operational complexity required, the possibility of increasing the energy by 5% really only applies when a machine already exists. We cannot guarantee this energy increase at the design stage.

For these reasons, using this technique to allow for a bigger spread in cavities is limited, so the primary goal of achieving a small spread in performance (high cavity yield) is still required.

Any other approach would require sorting cavities during production. This would be cost and schedule prohibitive, and would – in the end – not gain much more than playing with the LLRF.

2. Positron source:

a. What is the energy/orbit acceptance of the undulator for the 150 GeV beam?

In the design of the undulator carried out for the US Linear Collider Technology Option Study (USLCTOS), an 850 meter insert was used, which had space for twenty 10-meter undulator sections, sacrificial collimators, and diagnostics. The bandpass was quite tight (2%), primarily because of the strong focusing needed to limit the SR emittance growth – the matching to the linac probably could have been improved, although this was tried to some extent. The 2% bandpass corresponded a beam offset which exceeded the undulator aperture. The transverse offset of the insert was 2.5 meters to provide shielding.

It is certainly prudent to protect the undulator with a collimator (or collimators) and a fast extraction abort kicker system. In the present e+ layout, the collimators are upstream of the insert, as they should be to protect the system. The arcs could be made much shorter by placing the undulator at an angle to the main linac with a short arc at the beginning, and letting the photon beam drift. The electron beam could be bent back onto the original main linac axis, and the photon beam could drift for ~1 km to achieve a 3-meter offset that might be desired for shielding.

b. What kind of accuracy for energy/orbit control of the incoming beam is required?

From the point of view of positron yield, the allowed ‘energy jitter’ is probably significantly more than needed for other reasons (e.g, MPS, luminosity), and is not the driving constraint. The transverse beam jitter is a bit more stringent for the polarised source, due to the photon collimation. The photon beam divergence should be about five times the natural electron divergence to achieve the necessary polarisation. This would suggest that a 1 sigma electron beam jitter is probably negligible, and that one could even tolerate more. A fast feedback system upstream of the undulator could certainly achieve the necessary stability. For an unpolarised source (no photon collimation), probably no feedback would be needed.

c. What is the possible damage to the undulator by mishandling the incoming beam?

Clearly a direct hit of the full bunch train at 150 GeV would almost certainly destroy the undulator or a good part of it, as it would any other part of the machine. MPS must monitor the beam offset, and as mentioned above a fast abort system is prudent. Sacrificial protection collimators may also be needed.

d. What about the design of the collimation section to protect the undulator?

See comment above, part B2a.

e. If we do not have polarization at the beginning, can we reduce the length of undulator by assuming the acceptance of DR in Fig. 3 (p. 43)?

Yes, because you do not need to collimate the photon beam. But as polarised positrons are one of the attractive features of this source, you would still leave enough real estate for the ‘upgrade’, even if you started with an unpolarised source.

f. What about the depolarization from the target to DR?

The only ‘true’ depolarising effects should only be at low-energies, where space charge might play a role and the relative energy spread is large. If they have not already been done, simulations should be performed for the entire capture section (250 MeV point).

2. Comments on Damping Rings.

In his comments, K. Oide briefly re-stated the basic arguments for and against the 6 km rings solution (baseline) with respect to the dogbone. He notes that 6 km solution may be cheaper, although the current cost estimates are relatively crude. K.Oide’s comments on the fast kickers needed for the 6 km ring (particularly in conjunction with impedance and beam heating) should be addressed by the damping ring group. ATF can provide nearly the full-spec beam of the ILC damping ring, so a heating test can be done there when we create an ILC-like beam from ATF.

The basic issues with the choice of damping ring circumference come down to acceptance and space-charge, versus collective effects such as fast-ion or electron cloud. To these we should add commissioning and operational issues. Assuming that any dogbone solution adopted would share the linac housing, we must re-visit the question of the curved tunnel housing (this was looked at during the TDR stage, and the initial evidence suggests no real problems, but a more critical review would be warranted.)

In addressing the impact of the 5400-bunch low-Q parameter set on the 6 km ring baseline, the DR working group is beginning to study the limits on increasing the number of bunches beyond 2800, as the bunch charge is lowered. The GDE-EC encourages the DR group to carry out these studies as rapidly as possible.

In addition to the low-Q evaluation there are additional global questions which should be addressed. In addressing these issues, the DR group should work in concert with the GDE Accelerator Design Leaders to ensure that the solutions optimize the overall machine design. Two examples of such issues are:

a. Overall timing, bunch spacing constraints, and bunch train pattern flexibility. The damping ring is the heart of the machine in many respects; the DR parameters have many repercussions on the bunch train pattern and future flexibility.

b. Parallel operation of two stacked damping rings: this may result in additional constraints on overall timing, bunch spacing constraints, and bunch train patterns. If fast vertical kickers are required, then this may be a source of bunch-bunch random vertical jitter (but presumably this can be mitigated to some extent by feed-forward across the turn-around). If the rings are to be driven in parallel (half the bunch train to each ring), then the path length from each ring must be identical to within the phase stability requirement of the bunch compressor. This requires the two rings to have the same path length to within about 1 mm.

C. Selected points raised by Frascati Workshop Sessions, morning sessions: items to be included in the BCD

1. Main linac cryogenic system should be designed for 31.5 MV/m, $Q=10^{10}$ (with margin).

The BCD states that klystron power and cryoplant power must be installed for a gradient of 35 MV/m. The question was raised if this needs to have an additional overhead on the cryoplant power. If we assume a 50% overhead margin at the nominal running conditions quoted above, then at 35 MV/m (with $Q = 0.8x10^{10}$) , there is ~10% remaining overhead. This is acceptable, as the 35MV/m is in itself an overhead. It is not necessary to require a 50% overhead at 35MV/m (and it would be expensive).

2. The positron source undulator will use superconducting technology

One of the main arguments for a 'tuneable' S.C. magnet was that it could be turned off or adjusted, when the source was at the exit of the linac. For an undulator at the fixed 150 GeV energy point, this requirement is no longer applicable (although clearly it is still desirable).

3. The location of the keep-alive positron source should be specified.

There are two proposed solutions for the keep-alive source. Flöttmann has proposed an additional 500 MeV linac placed upstream of the photon target, which he estimates could give a 10% single-bunch charge (full bunch train). Apart from the linac, the source uses the same capture section and transfer line as the standard e+ source. Thus, it is necessarily on the electron linac side. Himel has proposed using a high-power target and capture section, which would make use of the 5 GeV positron injector linac to drive a conventional source with nominally half the current (probably less, when target constraints and damping ring acceptance are taken into account). This source will require a dedicated target, capture system, pre-acceleration and transfer line. If the e+ injector linac is located on the positron side, then this source has the advantage that it can run while accessing the electron linac or BDS.

Neither approach has been worked out enough to make a decision between the two, and given that the cost of either system is likely to be small, both should be taken further until a clear benefit of one source over the other is identified. However, the potential benefits of the Himel source in terms of decoupling the linac operations certainly demands that the e+ injector linac be placed on the positron side of the machine.

4.The location of the undulator: in-line, dog-leg, or chicane, should be specified.

The undulator insert should be as short as possible. The USLCTOS undulator insert had a length of about 850 meters, for a 2.5 meter offset. This length could be reduced with a non-symmetric chicane.

Allowing the linac to kink horizontally (a single bend solution) would probably be the shortest (and possibly the cheapest) solution. Next cheapest is probably a dogleg solution (half a chicane), with the full chicane solution being the longest.

The impact of our upgrade scenario needs to understood. In the upgraded machine, the undulator source will be at 650 GeV, unless we move it. This issue should be ignored in the 500 GeV machine design, since the disruption for the upgrade is now so great that ripping the source up and either moving it or reconfiguring it is not a serious impact. (although we should keep an eye on the required civil engineering for the source.) The exception to this is for the case of the kinked linac, which we can exclude from the start.

The dogleg solutions require a jog in the linac axis. To keep the same tunnel cross-section (layout), we will need to jog the tunnel also. The cost consequences of doing this need to be assessed.

The chicane solution (which preserves the main linac axis) would seem the most elegant, although at first glance may require the longest insertion (although a non-symmetric chicane would be shorter). This choice should be taken, but we should ask the e+ source groups to evaluate the impact of the other solutions.

5. The conflict between the GG2 and WG2 specification for the main linac BPM resolution should be resolved.

The current baseline BPM resolution requirement in the linac is ‘less than 10 μm’. This requirement comes from WG1, and is a direct result of simulations of dispersion-free steering. The WG1 members themselves are not in agreement here, with estimates ranging from a couple of microns up to the quoted 10 μm . Some simulations show that below ‘a few microns’ there is no more gain, since at that point presumably the wakefields from the misaligned cavities begin to dominate. Clearly there is much work for the simulators to do.

The specification from GG2 comes from a much more pragmatic (if not rather arbitrary) requirement, that the resolution should be a fraction of the beam size along the entire linac. This requirement is being driven by the need to analyze (and trace) sub-sigma beam jitter. Certainly the WG1 group never considered this.

More precise is always better, providing it doesn’t incur a significant cost increase. The requirement that all the BPM’s in the linac have sub-sigma resolution is not mandatory, although having some BPM’s capable of resolving sub-sigma jitter is clearly required (at fast-feedback locations, for example). However, it may be that there is no significant cost increase in achieving < 1 μm resolution. If that’s the case, then the instrumentation global group should certainly be encouraged to continue the R&D. Having such precise BPM’s opens up possibilities that we have not really considered, such as putting chirps or transverse signals in the bunch train during operations. Such ‘diagnostics’ may prove very powerful.

Specifying the <10 μm based only on simulations of DFS is perhaps weak, but at the same time, there is no justification for specifying that ≤ 1 μm is *required*. But it is certainly desirable if the cost is acceptable

6. The conflict between the GG2 and WG2 specification for the number of diagnostic sections in the main linac should be resolved.

From GG2 at Snowmass, there was a recommendation for an emittance measurement station at the entrance and exit of the main linac, and perhaps one at the 50 GeV point; the latter may be questionable. The proposal from GG2 for adding laser wire emittance stations to the cold machine without any significant increase in length certainly opens up the opportunity to add more. More were added by WG2 into the BCD. Also, GG2 did not recommend the need for any energy measurement diagnostics in the main linac (other than at the end), although again at least one has been added by WG2.

There is a danger in adding warm insertions to the cold linac, particularly when they at the same time introduce an energy bandwidth limitation. This opens up the potential for the beam to ‘punch a hole to air’ in the warm section and realize the nightmare scenario hinted at by B. Richter. We have to deal with this problem at the exit of the linac, and may consider adding some ‘beamline safety length’ before the first bend for just that reason. Adding additional such locations inside the main linac just increases the potential for problems.

A clear justification for the warm insertions added by WG2, other than those specified by GG2, should be provided.

7. Should the use of high-power processing for installed cavities be included in the baseline?

The GDE-EC requests that the main linac area managers and the SCRF cavity technical system managers come to a resolution regarding whether high-power processing for installed cavities should be added to the baseline.

D. Selected points raised by Frascati Workshop Sessions, afternoon sessions: discussion of white papers

1. The gamma-gamma option would be compatible with 14 mrad crossing angle and above.

GG6 should report on the compatibility of the single-IR configuration with the gamma-gamma option. The gamma-gamma option can have a strong impact on the choice of crossing angle. The GDE-EC requests the WWS and/or ILCSC provide the GDE with a statement of the priorities of the gamma-gamma option in the ILC physics program.

2. *Only the choice of the Marx modulator would allow the option of a single tunnel with the modulators in the tunnel.*

This issue has already been dealt with above, in part A6.

3. *There are ways to deal with the problem of He flow in the cryogenic system for a laser-straight tunnel, but they could be costly.*

There is no experience here, and certainly R&D is needed. In specifying the design of the cryogenic system for a laser-straight tunnel, it is important to appreciate the additional constraints of SC cavities (sensitivity to pressure, for example), which are not present in superconducting magnet systems.

E. Selected points raised by Frascati Workshop Sessions, afternoon sessions: general comments on the strawman BCD

1. Tom Markiewicz: Option 3 does not pass the "laugh test" for a 1 TeV upgrade.

2. Response from Nick Walker and Barry Barish: the options with an empty tunnel are not sellable as a baseline.

3. Jonathan Dorfan: With the Option 3 choice, we must now sell a 500 GeV machine, not a 1 TeV machine.

The physics case for the initial phase (500 GeV) is very strong. The case for a TeV machine is not so clear, and certainly would benefit from input from LHC. As was mentioned above in part A1, 1 TeV may just not be enough. Guidance must await LHC results.

4. *Hasan Padamsee: The cost differential of Options 1 and 3 is <15% of the TPC. If cost is the main driver, it should also drive the choice of 1 tunnel over 2 tunnels.*

The current increase in cost for two-tunnels can be justified.

The argument for Option 3 is not simply a cost issue, but also 'sellability' issue. The project must be "sellable", and manifestly cost-optimized, not just to funding agencies, but to legislative bodies, scientists in other fields, and the general public.

5. Eckhard Elsen: There is a good case for 500 GeV; we can wait until the LHC comes on to determine the upgrade energy (if any).

This point is made above. One would hope that the LHC (or initial ILC results) do point to an upgrade energy.

6. *PT: The gradient goals are challenging; the choice of a fixed length tunnel puts the final energy at risk.*

There is a inherent and real risk in achieving the maximum energy irrespective of what we specify. We will always have a fixed tunnel length, and the top end of the machine will be dictated by the performance of the hardware installed.

If routine operation at 500GeV is important (physics driven?), then we must certainly build enough overhead in to guarantee that we achieve it, based on our best (and conservative) estimate of the expected performance. We either build more linac from the start, or we have some amount of free space (empty tunnel) to allow us to add more linac if needed (as PT implicitly suggests). Either way we will need to justify the overhead and associated cost, which ultimately means justifying 500 GeV (or what ever your favourite number is), irrespective if you call this 'overhead' or 'upgrade'.

7. *Ron Settles: Physics will determine what upgrade is needed and provide the justification for it.*

A strong case made today and independent of LHC would certainly be very compelling.

8. *Nobu Toge: Agrees with the choice of option 3. But the GDE EC should write down the "sellability" argument for option 3.*

The GDE-EC will provide text to make this argument, in the BCD (under justification for the energy upgrade).

9. *Francois Richard: We could sell the TESLA TDR as "full coverage" of the physics, but the limit to 500 GeV requires a light Higgs. But can we react to LHC?*

While discovery of a light-Higgs could almost be interpreted as a blank check for a 500 GeV linear collider, it is not by any means the only argument. There are strong physics arguments in the no-Higgs scenario also, not least that LHC may have missed it.

10. Bob Kephart: Cryomodule fabrication capability will need to be re-started long after the initial stage of the project.

True, but this is more of an argument to maintain the initial rate and construct the entire 1 TeV machine in one go. Dropping down the production rate to a 'trickle' over (say) ten years to adiabatically upgrade the energy would also need to be factored into the cost estimates and agreements with industry from the start. The funding agencies will be aware of this, and may argue that the entire upgrade to 1 TeV must be justified initially.

11. *Nigel Lockeyer: The 500 GeV physics case is compelling: there will be a scalar fundamental particle to study.*

12. *Phillipe Bambade: Are we abandoning 1 TeV completely?*

13. Response from Barry Barish, Nick Walker, Ewan Paterson: No, only the tunnel. All else required: land acquisition, dumps, etc, will be appropriate for 1 TeV.

No other HEP experimental facility (except LEP) has ever been required to have the potential for a factor of 2 or more increase in centre-of-mass energy. We can always build more linac. There is actually no real reason not to go higher than 1 TeV, providing that space, AC power and money are available.

14. *Tor Raubenheimer: Using change control, we can resize the machine when more information is available.*

15. Jonathan Dorfan: Change control notwithstanding, what we are selling is less justifiable (for physics) than project with long tunnel.

One can again turn this argument around (as pointed out above), that we cannot justify 1 TeV because we do not know enough.

16. Tom Himel: The Marx modulator should be the baseline, we are being too conservative. Changing to the Marx will require many changes in tunnel layout, electrical systems, etc. Better to do these now.

17. Marc Ross: The Marx modulator is still a ways from completion; there are still serious challenges.

18. Nigel Lockeyer: When will Marx be tested?

19. Ray Larsen: IGBT's will be delayed, but still on schedule for summer test. Recommends sticking with the Bouncer modulator for now.

20. Tor Raubenheimer: Probability of success of Marx modulator is high. But this change could be deferred to CCB process (but soon).

This will get resolved one way or the other when the (inevitable) change request is made. At that point it must be reviewed.

21. *Eckhard Elsen: Role of modulator is important. What is the coupling to single tunnel? Should number of tunnels be reconsidered if Marx is chosen?*

See above, part A6.

22. *There should be a uniform way to "sell" the project with regard to costs. Apply same cost-sensitive strategy to all parts of the project.*

This is certainly true. However, the energy upgrade option choice involves more than simply costs, as discussed above.

23. *Peter Garbincius: We are optimizing to a cost number that we don't know. We should adopt a strategy to get costs early (June) rather than late.*

24. Nick Walker: We should have a ball-park number by the Vancouver meeting.

The ball-park number must also be believable. We need a clear strategy to specify how we will determine the main linac costs for the RDR.

25. *Nan Phinney: Regarding the curved tunnel: what is the cost savings for the cryogenics?*

R&D is needed. At CERN, the laser-straight solution for their site is believed to be more cost effective from civil engineering point of view; they have not considered the cryogenics yet.

26. *Jonathan Dorfan: Will the final cost be a number or a range? How will the costs be quoted?*

27. Bob Kephart: Do we really need to quote a cost? Or only a partial, regional cost, as in ITER?

28. Barry Barish: We will have to quote a vetted cost. We will need to translate our probabilistic risk approach into a range for the cost.

29. *Gerry Dugan: What is the number that we can't exceed?*

30. Paul Grannis: We won't know what the number is.

31. *Barry Barish: We will have only one chance to sell the project: we just have to come in with the best case we can make.*

32. *Nobu Toge: Analogy with TRISTAN: shoud we try to sell a higher energy machine, so that there is some room to back off when less is supported? How do we interact with Asian and European funding agencies?*

33. Jean-Pierre Delahaye: We will need to fit the funding for the project within the existing global HEP budget profile.

F. Comments received by email on the strawman BCD

1. Weiren Chou: One of the cavity gradient selection criteria is to minimize the cost. But I found the plot showing ILC cost vs. gradient was confusing. The civil and cryomodule (both ~1/G) cost much more than the cryoplant (~G^2). The former is more than 40% and the latter only about 4% of the total cost. How can the curve look symmetric near 40 MV/m?

The cost fractions that you cite correspond to a gradient choice which is well below 40 MV/m, the minimum cost point. For this gradient choice, the cryogenic costs, which vary with G^2 as you note, are much smaller than the other costs. As one increases the gradient to approach 40 MV/m, the cryogenic costs rapidly increase, while the civil and cryogenic costs decrease, resulting in the curve shown. The curve is not really symmetric, but tends to rise more sharply on the high gradient side, because the cavity Q decreases with gradient at high gradient, resulting in cryogenic costs which rise more rapidly that G^2 at high gradients.

2. Marc Ross, Junji Urakawa, Hans Braun (GG2 convenors):
Concerning: BCD Linac beam position monitor performance criteria

The 'strawman BCD', as posted November 21, lists the linac BPM performance requirement as '10 microns or better'. By this letter, we request that this requirement be discussed and revised.

At the Snowmass meeting, Global Group 2, Instrumentation and Controls, recommended that the linac beam position monitor performance should be based on operational issues, namely that the required resolution should be a fraction of the beam size (much less than 10 um.). This topic was discussed and agreed upon at our joint meeting with WG1. The group as a whole was in agreement on this point and made sure the recommendation was clearly explained and widely published (posted November 6).

It is our understanding that the development of the BCD was to be based on the recommendations made at Snowmass. The WG2 BPM performance statement was made without any discussion with the group GG2.

The group (GG2) would like to ask that this discussion take place so we can understand the reason our recommendation was ignored.

See the response documented in Section C5 above.

3. Amin Reichold: I noticed that the current strawman does not contain a section for the survey and alignment of the ILC. At Snowmass this subject was quite successfully discussed iWG1 (LET) and GG5 (instrumentation). I would be glad if you could let us know where you wish this section to appear in the BCD.

The baseline method for survey and alignment of the machine components should be specified in the BCD. The GDE-EC recommends that the CCB determine the correct section of the BCD for this to appear in, and work with WG1 and GG5 to document the baseline method, as determined at Snowmass.

4. V.I. Telnov: I have refreshed on Nov.23 the GG6 contribution . New is the figure and some discussion of possible configuration of the IR with a large crossing angle.
Three schemes are considered. In the first the crossing angle for gg and e+e- is the same, about 25 mrad. In the 2-th and 3-th the e+e- crossing angle is 20 mrad or somewhat less. Note that, if the beam dump for e+e- and gg is the same (which is desirable), it should be not too far (250 m or so), otherwise the radius of the disrupted beam at gg,ge will be too large.
It is necessary to consider GG6 suggestions by WG4, make some relevant consideration on required upgrade from e+e- to gg,ge in the "beam delivery" chapter and reflect this in BCD.

See Section D1. The GDE-EC agrees that the requirements of the gamma-gamma option should be considered when determining and evaluating the baseline choice for the IR configuration.

5. Gudrid Moortgat-Pick: This comment was a 3-page letter. Included here is only the summary.

Summary:
It is of utmost importance that the BCD is cost optimized for the 500 GeV energy. The TeV upgrade considerations should not cause major cost-driving design specifications for the 500 GeV option. A clear physics case for ILC(500) exists, based on top physics, Higgs physics and a high potential for detecting effects of new physics both in direct as well as indirect searches. Early results from the LHC will most likely provide further support for this energy stage, so that an expeditious realization of the ILC(500) should be pursued.

Highly desirable features of ILC(500) are

a) a tunable energy to lower energies,

b) tunable energy to higher energies beyond 500 GeV at cost of reduced luminosity,

c) the polarization of both beams at a very early stage.

These features would allow to maximally exploit that energy stage, provide a *promising 'safety margin' for the physics potential and allow to determine –together with the LHC results– the most suitable upgrade strategies.*

See the responses above to comments A1 and E1-E14. The GDE-EC agrees that the machine should be cost-optimized for 500 GeV, while not excluding extension to 1 TeV at a future time. This is one of the reasons for the choice of "Option 3" for the energy upgrade path. The machine must also be easily tuneable over a wide range of energies. This is one of the reasons for the choice of the undulator location at the 150 GeV point. The choice of an undulator source also allows a simple and relatively inexpensive upgrade path to a polarized positron beam.

Regarding the choice of gradient, "Option 2", with a lower initial operating gradient (28 MV/m), was not chosen primarily for reasons of cost. The baseline option, with a gradient of 31.5 MV/m, but with cavities capable of operation up to 35 MV/m, will in principle allow the machine to reach an energy of about 550 GeV, with reduced luminosity. While this energy reach is not as great as that of Option 2, the added cost of Option 2 is significant and was not felt to be justifiable. As pointed out above, the physics case for 500 GeV is strong, while extensions of the machine above this energy in the initial stage cannot be strongly motivated.