U.S. Linear Collider Technology Options Study

U.S. Linear Collider Steering Group
Accelerator Sub-committee
Linear Collider Option Task Forces

March 4, 2004

Contents

List of Figures

List of Tables

Chapter 4

Availability Design

4.1 Introduction

As part of the Technology Comparison exercise, an availability task force was formed to assess the impact of various technology or configuration choices on overall machine performance. The approach taken was to set an overall availability goal for any linear collider and then use that to develop an unavailability budget, which apportioned the down time among the different systems of the accelerator. The required availability specifications were compared to the performance of existing accelerators to see how much the reliability of individual components must be improved. This level of improvement could then be used to determine the added costs needed to achieve the budgeted availability and to assess the risks of not achieving the necessary improvements.

Note that the goal was **not** to determine the availability of each proposed option or variant and declare it a bad design if it was down too much. Rather, the assumption is that, with proper design, either of the options or any of the variants can be made available enough; it is only a matter of how much it will cost, and the risk of failure if components must be made too much more reliable than the present state-of-the-art.

Given the limited time available, this study is of necessity only a first crude step. Nevertheless, it has produced a very useful tool to evaluate the overall availability of an accelerator. While the tool has not been benchmarked against an existing accelerator and overall predicted downtimes could change with different assumptions, we believe the comparisons between different accelerator variants (warm/cold, 1 vs. 2 tunnel, conventional vs. undulator e^+ source) are quite significant.

Before proceeding, it is worthwhile to give a few definitions.

availability : the fraction of time the accelerator hardware is up enough to be providing useful beams to the detector. For this exercise, substandard performance while tuning emittance or backgrounds, or recovering from MPS trips is still considered to be available. The luminosity loss due to these items must be accounted separately.

unavailability : just one minus the availability. For this exercise, time spent tuning the accelerator to recover after a repair downtime is considered as unavailable.

reliability : sort of a synonym for availability, not used in a precise manner.

MTBF : the Mean Time Between Failures is the average time between failures of a component assuming normal preventive maintenance has been done. It is very important to realize that if a component is specified as having an MTBF of 50 million hours that does **not** mean it must run for that long

without being touched by human hands. Fans, bearings, or even the whole device can be replaced on a regular basis. Only breakages of the device which occur outside our control count toward the MTBF. The MTBF as used here does **not** include infant mortality as it is the steady state running of the accelerator that is being considered; hence the devices have been burned in. MTBF is **not** not the lifetime of a device; it can be much longer than that since one can use preventive maintenance to replace the wearing part or the whole device before it wears out. It is only if one neglects to do this preventive maintenance that the MTBF will be less than the lifetime.

MTTR : the Mean Time To Repair a component. This is the time from when one has access to the component until it is repaired. It includes travel time, but not time for the accelerator to cool down before an access starts, nor time to turn on and tune up the accelerator.

region : a major part of the accelerator such as the e^- injector, e^+ damping ring, or e^- linac.

MD : Machine Development or time devoted to accelerator physics experiments, including trying out possible improvements.

opportunistic MD : Machine Development done in a region of the accelerator while some other region(s) are undergoing repairs or are still recovering from a repair.

4.2 Method

A detailed list of components was compiled for two major regions of the collider, the linacs and damping rings (DRs). These regions were chosen for detailed analysis both because of their size and complexity, and because they were the areas that change most significantly depending on the main linac technology. The component list included items identified as potential sources of failure from experience with existing facilities, such as RF components, magnets, magnet power supplies, power supply controllers, vacuum pumps, pump power supplies, movers, diagnostics, etc., with counts for each item.

Each component was then assigned a starting MTBF and MTTR. Where possible, the numbers were based on data from repair statistics at SLAC, Fermilab or other labs. In a few cases (e.g. magnet power supplies), the linear collider proposals had already engineered a level of redundancy to improve the MTBF. Here, the starting numbers assumed the designed redundancy and those cases are explicitly mentioned in the results section. In some cases, no good data was available and the MTBF used was just a reasonable guess (typically 100,000 hours). It is worth noting that even for the most common devices, MTBFs can vary by a factor of 10 depending on the exact design and application of the device.

Due to lack of time, the other regions of the accelerator were not modeled in detail. Rather, the sources and beam delivery systems (which are very similar in all designs) and bunch compressors were treated as individual elements with an MTBF and MTTR for the whole region.

To develop a more quantitative evaluation, a simulation was written which used all of this input data to estimate the availability of the whole LC. A brief description of this simulation is given in the next section, and more complete details are available in LCC note 127.

The simulation assumes the accelerator has reached a steady state after several years of operation. It does not use a bathtub curve which would give poorer MTBFs in the early years due to infant mortality and in later years due to devices wearing out. A special run was made with all parameters degraded to simulate what operation might be like in the early years.

The results of the simulation gave the total unavailability and how each device type contributed to it. As expected, the availability of the LC was too poor and it was necessary to improve the MTBFs of some devices to attain the goal availability. Typically, the devices which were contributing more than most to

the unavailability were identified and improved. The rationale was that it would be cheaper to concentrate efforts on a few critical devices rather than trying to improve all devices by the same factor.

There is clearly flexibility in which devices to improve in order to achieve the goal availability. Ideally the improvements chosen would be those which minimize cost increases. While there was not time to do this quantitatively, this report provides a reasonable first iteration. As a design progresses more, these tools should clearly continue to be used to optimize the design.

The final result is a table of the MTBFs of devices that are needed to achieve the desired availability with an indication of how much these must be improved over the present state of the art. There are also comparisons of how the availability would change when going from 2 tunnels to 1 tunnel or from an undulator positron source to a conventional source.

4.3 Brief description of the simulation

Many accelerators have estimated their availabilities during the design phase with a spreadsheet. (Examples from SNS and APT were reviewed.) There are formulas to combine the availabilities of the components to get the availability of the whole. There are also reliability software packages that help do this. The approach taken here was to write a simulation which could allow several complexities to be handled that would have been nearly impossible in a spreadsheet and quite difficult in the commercial software packages. These complexities include the recovery and tuning time needed after a downtime, the complex redundancies built into the LC designs, the way in which accelerator physics experiments (Machine Development or MD) can be done when only part of the accelerator is down, and the way in which many devices are typically repaired during an access by a limited number of people. By writing a simulation tailored to the task, it was possible to incorporate the decades of experience in running real accelerators accumulated by the people on the task force. The following paragraphs give a brief description of these aspects of the simulation. LCC note 127 gives details for those who are interested.

4.3.1 Handling Complex Redundancies

If all the devices in the linac had to be working at the same time, the LC would never be up. All designs have a built in energy overhead so that some energy producing components can be broken and the LC still functions because an energy feedback loop keeps the energy constant. Many components such as the klystrons, modulators, RF structures, AC power for the RF, water pumps, and tuners will degrade the energy when they malfunction. Some of these can be repaired while the accelerator is running. Others must await an access. Each degrades the energy by a certain number of MeV. The simulation starts with a perfect accelerator with the energy overhead given in the design. Each time a component breaks, it decreases the energy overhead by the corresponding amount and schedules it for repair (either immediately if it can be done "hot" or at the next downtime). When the repair is completed, the energy overhead is increased. If the energy overhead gets down to zero, the accelerator is declared broken and many accumulated repairs get done.

The two linac (e^+ and e^-) energy overheads are individually tracked. Each warm DR RF system is also redundant (with one more cavity or RF station than needed) and is handled in a similar fashion. Also the cold DR's injection and extraction kickers are redundant (21 kicker systems where only 20 are needed).

Things also break which can degrade the luminosity without turning off the machine. Examples are diagnostics such as BPMs and laser wires which make tuning more difficult and hence reduce the luminosity. For other devices such as such as correctors or movers in both warm and cold LCs or linac quadrupoles in the cold LC, it is assumed that after some downtime for retuning, the machine can continue to run with the broken device. The simulation makes each of these broken devices degrade the luminosity a bit (0.1%

per BPM and 5% per wire). If the luminosity gets down to half its design value, the accelerator is declared broken and many accumulated repairs get done.

If an item that simply degrades a parameter can be repaired while the accelerator is running, the simulation does so, as long as sufficient repair people are available. These are called hot repairs. The repair of klystrons and modulators are the most important examples of this.

Of course, some components are not redundant and when they fail, the accelerator is considered broken and repairs must be scheduled immediately. Most magnets other than correctors, together with their power supplies and controllers, are typical examples. In the cold linac, however, the quadrupoles are treated as a special case. The focusing lattice is weak enough that the simulation assumes that one can tune around the problem in a couple of hours to recover 99% of the luminosity. The failed component is then repaired at the next convenient time (typically during the long down for the magnet which is in a cryostat). Other devices handled in a similar fashion are movers in the warm linac and DR, trim windings and correctors in both DRs, and tuners and couplers in the cold linac.

4.3.2 Downtime planning

Without doubt the downtime planning is the most complicated part of the simulation. This should come as no surprise to anyone who has participated in the planning of a repair day. It is even harder in the simulation because computers don't get a gestalt of the situation like humans do. Briefly, it figures out what parameter (e.g. e^- linac energy overhead or e^+ DR extraction kicker strength or luminosity) got degraded too much, and plans to fix things that degrade that parameter. Having figured out how long the downtime must be to fix items which simply have to be repaired, it then schedules other items for repair allowing them to extend the downtime by 50 to 100%. Some other issues must also be taken into account:

- If an access to the accelerator tunnel is required, 1 hour is allowed for prompt radiation to decay before entry. 1 hour is also allowed for locking up, turning on and standardizing power supplies. (This cooldown time is what is typically required at SLAC to minimize exposure of employees to residual radiation.)
- The devices chosen for repair are those that give the most bang for the buck (most improvement in the parameter per hour of repair time).
- The number of people in the accelerator tunnel is limited to 50 (to minimize the chaos of PPS control). This number is consistent with practice at existing facilities. The number of people doing other repairs is limited to 100 to reflect the finite resources available. All repair people are considered equal and specialties are not tracked. It will be shown later that the number of personnel could be limited still further without degrading the availability.
- There are no regularly scheduled maintenance shutdowns, except yearly. Interventions occur only when the accelerator is broken, which is what happens at most operating accelerators. In real life, maintenance might be planned when the energy overhead was getting low without waiting to actually run out of energy. However, since the simulation does not add any penalty for unplanned or off-hours downtimes, this becomes a subtlety which does not really impact the results.
- Things which break during the downtime are just ignored (assumed they are immediately fixed). The long recovery time which is described in the next section is intended to account for this.

4.3.3 Recovery times

The simulation assumes that all repairs get done on schedule. It seemed an unnecessary complication to throw random numbers to distribute the repair times around the MTTR as the simulation is run for long

enough to average out such variations anyway.

Recovery of the beam is modeled in a crude fashion which matches the qualitative experience on many accelerators. This common experience is that it takes time to recover good beams after a downtime. In fact, the longer the down, the longer the recovery time. Contributions to the recovery come from myriad factors such as

Hardware failures - devices such as pumps and power supplies which break because they were turned off during the shutdown or devices which just happen to break while the accelerator was down and were not detected

Environmental factors - temperature changes caused by the access or ground motion over a few hour period which can be significant enough to require retuning

Human error - mistakes made in doing the repairs (valves left closed, cables left disconnected...) or failure to restore settings after hardware or software tests

Parameter drifts - multiple parameters which are continuously tracked and optimized during normal operation which all need to be identified and retuned

Commissioning - hardware or software improvements made during the shutdown which need to be tested, calibrated, etc.

...

Rather than trying to model recovery procedures in detail, the simulation simply assumes that the time it takes to get good beam out of a region of the accelerator is proportional to the time that region was without beam. Table 4.3.3.1 shows the constants of proportionality used for each region (typically 10%, except for the DRs and IR, which are given 20%). In real operation, the beam quality recovers gradually as each region is tuned up in succession, and the luminosity gradually ramps up to nominal. The simulation simplifies this by assuming that the machine goes from no beam at the end of a region to perfect beam at the end of the recovery time. Similarly, the luminosity jumps from zero up to the design value immediately at the end of the recovery/tuning time. While this is certainly an oversimplification, if the recovery time used in the simulation is considered to be the time it takes to get back to half the design luminosity, then the overall effect is reasonably well reproduced.

4.3.4 Lumped Systems

As mentioned in Section 4.2, some systems were not modeled in detail due to lack of time, but rather modeled as one piece. The details are given in Table 4.3.4.1. Each such system was considered to have two types of failures: one which required access to the accelerator tunnel (flagged by a "1" in the "access needed?" column) and one which did not require access. A nominal downtime was chosen for each system for each access state, typically 0.2%. The warm bunch compressors have two stages and are more complex so they were given 0.3%, and the somewhat simpler IP regions and cold bunch compressors were given 0.15%. These downtimes were then converted into effective MTBFs. The output of the simulation shows a downtime for these systems which is roughly double the nominal downtimes in the table due to the modeled recovery time. Together these lumped systems required two-thirds of the unavailability budget and the rest was allocated to the linacs and DRs.

4.3.5 Machine Development

Machine Development is an essential tax on the operating efficiency of any accelerator. It is time used to better characterize the machine, develop new tuning procedures, and test possible future improvements. The

Table 4.3.3.1: Regions of the machine with allocated tuning time and MD. The first column gives the names of all the regions. The second gives the ratio of the time needed to recover beam to the time beam has been absent from the region. The third column gives the percent of scheduled running time that is needed to do MD for the warm design. Column 4 gives the same for the cold design. Note that the MD time includes the time to recover from the MD work. Hence the useful time to do actual MD measurements will be somewhat less than the times given in the table.

Region name	Tune time fraction	% MD time warm	% MD time cold
e^- injector	0.1	1	1
e^- DR	0.2	2	2
e^- compressor	0.1	1	1
e^- linac	0.1	1	1
e^- BDS	0.1	1	1
e^+ source	0.1	1	1
e^+ PDR	0.2	1	0
e^+ DR	0.2	2	2
e^+ compressor	0.1	1	1
e^+ linac	0.1	1	1
e^+ BDS	0.1	1	1
IP	0.2	1	1

Table 4.3.4.1: Details on how the lumped systems were modeled.

System	access needed?	MTTR (hrs)	warm nominal % downtime	cold nominal % downtime	warm MTBF (hrs)	cold MTBF (hrs)
e^- injector	1	8	0.2	0.2	4000	4000
e^- injector	0	1	0.2	0.2	500	500
e^- compressor	1	8	0.3	0.15	2667	5333
e^- compressor	0	1	0.3	0.15	333	667
e^- BDS	1	8	0.2	0.2	4000	4000
e^- BDS	0	1	0.2	0.2	500	500
e^+ source	1	8	0.2	0.2	4000	4000
e^+ source	0	1	0.2	0.2	500	500
e^+ PDR	1	8	0.2	1×10^{-40}	4000	8×10^{42}
e^+ PDR	0	1	0.2	1×10^{-40}	500	1×10^{42}
e^+ compressor	1	8	0.3	0.15	2667	5333
e^+ compressor	0	1	0.3	0.15	333	667
e^+ BDS	1	8	0.2	0.2	4000	4000
e^+ BDS	0	1	0.2	0.2	500	500
IP	1	8	0.15	0.15	5333	5333
IP	0	1	0.15	0.15	667	667
Cryo plants	0	10	1×10^{-40}	1	1×10^{43}	1000
site power	0	10	0.5	0.5	2000	2000
global controls	0	5	0.2	0.2	2500	2500

amount of time spent on MD varies through the life of a project, with more MD required in the early stages or after a major upgrade. For this simulation, the LC is assumed to have operated for a few years and to have settled into a nominal schedule of MD, which would occupy approximately 10% of the time. As with the recovery time, the MD was allocated to the individual regions of the machine. Each region was allocated 1% MD with the exception of the DRs, which were given 2% because of the complexity of the invasive tuning anticipated.

To more fully mirror the complexity of operation of a real machine, the simulation assumed that some of the required MD could be done on an opportunistic basis. Typically repairs may be completed in some regions earlier than others. As an example, there could be 16 hours of repair work needed in the e^- linac, 2 hours in the e^- injector and none anywhere else. In this case, there will be time after the e^- injector is repaired when beam has been tuned up into the e^- DR while repairs are still going on in the e^- linac. If this is more than two hours, the simulation assumes useful MD can be done in the e^- injector and DR. If there is a conventional positron target, MD can also be done anywhere in the e^+ system. This opportunistic MD time is tracked by the simulation. It then assumes that sometime during the running period enough scheduled MD is done in each region to bring the total of opportunistic plus scheduled MD up to the desired levels given in Table 4.3.3.1. Note that the MD times in the table include recovery from the MD. During scheduled MD the LC is up, but not producing luminosity for the detector.

4.3.6 Important assumptions that effect availability

There are several assumptions incorporated into the simulation that are important and are not explicitly stated in the description of the accelerators. If the final LC designers are foolhardy enough to not meet these design requirements, they should be sure to repeat the simulation to see the effects of their less reliable design.

1. Klystrons located in the support tunnel can be replaced while the accelerator is running. This requires a valve or extra window in the high power waveguide, which is not implemented in the existing test facilities, TTF or NLCTA.
2. All electronics modules and power supplies are accessible in the support tunnel including vacuum pump power supplies and BPM processors. Electronics modules are hot swappable; that is, they can be replaced without turning off other electronics.
3. There is a tune-up dump at the end of each PPS zone so that beam can be run in an upstream region while people are in the tunnel performing repairs downstream.

4.3.7 The need for bench marking (and why it is hard)

During the development of this simulation, progress reports were given to several groups. While most people agree that the technique looks good, someone always asks, "Shouldn't you benchmark this against an existing accelerator to make sure it is working properly." The answer is "Yes, but." Yes, it should be done, but it is a lot of work and there was not enough time to do it. Hopefully, future work with the simulation will include benchmarking it against a real accelerator. This is a nontrivial task for several reasons:

- Getting together the list of components is real work
- MTBFs and MTTRs for that actual accelerator should be used if one wants an accurate comparison. Again, it is real work to get these from whatever trouble reporting database the accelerator keeps. For the cases studied (SLC and Fermilab), there were often more errors than correct entries in the trouble database and it was necessary to go through thousands of errors by hand and correct the entries.
- Most labs record the recovery from a downtime as "tuning", not as downtime attributed to the device which originally broke. One would either have to modify the simulation to reproduce what the lab did (trivial, but misleading), or go back through a few years of logbooks and try to correct the data for the accelerator.
- Most labs try to schedule repairs before they completely bring the accelerator down (e.g. a water leak that hasn't yet caused a magnet ground fault or the repair of a water pump whose bearing is screeching)

at least a day in advance. They then account this as scheduled downtime (no different than a scheduled Christmas shutdown or 3 month shutdown) rather than unscheduled downtime. Whether it is scheduled or unscheduled, no luminosity is made, so the simulation just calls it downtime. The data from a real accelerator would have to be corrected for this effect.

4.4 Results

The task was to create an unavailability budget for several variants of the LC. It was decided to allow 25% downtime due to hardware problems and the recovery from them. This is comparable to what present HEP accelerators accomplish after several years of effort. The LC is much more complex, so it will be more difficult to achieve. Of the 25% downtime, only 15% was explicitly budgeted to specific devices. The other 10% was held as contingency for devices that were not included or for major design mistakes that cause poor reliability and that remain a problem after a few years of operation. To make it crystal clear what is and is not included in the 25%, it is worthwhile to list all of the factors used to turn peak luminosity into integrated luminosity.

HA is the Hardware Availability. This includes the recovery from a hardware problem. This is the only thing has been studied in detail with the simulation. A 25% downtime is allowed or 75% uptime.

MD is the fraction of time spent doing scheduled Machine Development. If there were no opportunistic MD done, then scheduled MD would be about 7% for the conventional e^+ source and 12% for the undulator source. The actual number varies from 2% to 11% depending on the accelerator variant. This variation is due to the use of opportunistic MD and to the possibility of doing MD in two parts of the accelerator at once. The simulation produces an estimate of MD, but for this example, an average of 10% will be used.

SD is the fraction of time in the long Shut Down. A 3 month shutdown once every year gives 25% for this.

SU is the fraction of time Starting Up and Recovering from the long shutdown. Typically the luminosity ramps up gradually to the nominal value. Consider SU to be the fraction of a running year to get to half the nominal luminosity and then if the ramp were linear (which it usually isn't) then the fractional loss in luminosity is simply SU. For this example a 1.5 month recovery from the 3 month shutdown will be used giving SU = 1.5/9 = 16.7%

MPS is the fraction of time lost to MPS trips (and recovery from them) and other similar very short outages. For this example 5% will be used.

DT is the De-Tuning factor. It is the fraction by which the average non-zero luminosity is lower than the peak luminosity. Contributors to this are non-optimum tuning due either to mistakes or to the accelerator drifting away faster than feedbacks and operators can tune. For this example 10% will be used.

Taking all this into account then gives:
Integrated Luminosity per year

$$= \text{Peak Luminosity} \times \text{Seconds/year} \times \text{HA} \times (1\text{-MD}) \times (1\text{-SD}) \times (1\text{-SU}) \times (1\text{-MPS}) \times (1\text{-DT})$$
$$= 2 \times 10^{34}\text{cm}^{-2}\text{s}^{-1} \times 3.14 \times 10^{7}\ \text{s} \times 0.75 \times 0.90 \times 0.75 \times 0.833 \times 0.95 \times 0.90$$
$$= 2.26 \times 10^{41}\ \text{cm}^{-2} = 226 \text{ inverse femtobarns per year}$$

4.4.1 Simulations of the warm and cold reference designs, the cold 1 tunnel variant, and the conventional e^+ variant

Table 4.4.1.1 shows the results of simulation runs for the warm and cold reference designs, the cold 1 tunnel variant, and the conventional e^+ variant. Each run is given a unique name such as Warm1 or Cold3 in the first column and this name is used here to identify the run as each is described in turn.

Runs Warm1 and Cold1 simulate the main warm and cold LC options which have 2 tunnels (one for the accelerator and a second for support equipment) and an undulator based positron source. Nominal MTBFs (taken from operating accelerators or educated guesses when that data was unavailable) were used with the following exceptions:

1. All magnet supplies for both warm and cold use an MTBF of 200,000 hours instead of the 40-50 thousand hours experienced at SLAC and Fermilab. This comes from a design done for the TESLA TDR which has redundant regulators (which are the most common component to fail in a power supply) and hence an MTBF that depends mostly on the low power components and bulk supply which were not made redundant.

2. The power coupler interlock electronics and sensors for the cold LC were given an MTBF of 1 million hours instead of the nominal 100,000 hour MTBF that is used for most electronic modules and sensors. This was done because it is planned to build redundancy into the circuits and sensors.

3. The cold design calls for the cavity tuner stepping motors to be inside the cryostat. Using a nominal MTBF of 500,000 hours (taken from SLAC mover experience) resulted in so many tuner failures that a multi-month access was needed to perform repairs. To keep this from skewing the results, 1 million hours was used. It will become evident later that it will be necessary to make these components redundant or to bring them outside the cryostat.

4. The cryo plant for the cold linac is assumed to have a 99% uptime. As it is made of six systems each with a capacity of 30 kW at 4 K, each system must have an uptime of 99.84%. Each system is comparable in size to the Fermilab central helium liquefier which after recent major improvements is up about 99.6% of the time (not counting outages due to utilities like power and cooling water). CERN's LEP refrigerators were about half that capacity and were up about 99.83% of the time (again not including utility outages). Including utility outages tends to double or triple the downtime. Depending on what you compare to and whether outages due to utilities are included, the cryo plant's MTBF is a factor of 1 to 6 better than present experience.

5. The MTBF of present SLAC modulators is 30,000 hours with many of the failures coming from the thyratrons. The LC modulator designs all use IGBTs instead and hence an MTBF of 50,000 hours was used for them. In addition, the modulator design for the warm linac has built in redundancy as there are extra IGBT boards which allow the modulator to deliver full performance even with failed boards. For this reason it was assigned an MTBF of 100,000 hours.

A description of all the devices and the MTBFs and MTTRs used for them may be found in the appendix. The actual input and results spreadsheets and the simulation code used in all the runs may be found on the website: www.slac.stanford.edu/∼tmh/availability/. With only the above exceptions to nominal present day MTBFs, the downtimes of the warm and cold LCs end up being 28 and 31% respectively (taken from the number in the third column of the table). These are remarkably similar, and remarkably good given how complex the LCs are. However, it will not be trivial to achieve the goal of 15% down. The regions not modeled in detail (injector, bunch compressors, BDS, site power, cryo plant, global control system) were given fixed downtime budgets to cause a downtime of about 10%. Hence the linac and DRs which are modeled in detail cause about $30 - 10 = 20\%$ downtime with the nominal MTBFs. This must be reduced to 5% or a factor of about 4 to fit into the budget.

Table 4.4.1.1: Summary of simulation results. Each line represents one run of the simulation. The 10% downtime that is held as contingency was not simulated and hence is only reflected in the table in the column marked, incl. 10% contingency. Details of the MTBFs used in versions A, B, and C are given in Tables 4.4.1.2, 4.4.1.3, and 4.4.1.4. Note that columns 3 and 4 should add to 100% as should columns 6-9.

Run number	LC type	Simulated % time down incl. forced MD	Simulated % time fully up integrating lum or sched MD	Total % time integrating lum, incl 10% contingency	Simulated % time integrating lum	Simulated % time scheduled MD	Simulated % time actual opportunistic MD	Simulated % time useless down	Simulated # of accesses per month
Warm1	Warm, 2 tunnel, nominal MTBFs, und e^+	28.1	71.9	54.7	64.7	7.2	5.8	22.3	6.0
Warm2	Warm, 2 tunnel, vers A MTBFs, und e^+	15.0	85.0	64.6	74.6	10.4	2.6	12.4	3.3
Warm3	Warm, 2 tunnel, vers A MTBFs, conv e^+	11.3	88.7	77.7	87.7	1.0	7.0	4.3	2.8
Cold1	Cold, 2 tunnel, nominal MTBFs, und e^+	31.5	68.5	49.2	59.2	9.4	2.6	28.8	11.1
Cold2	Cold, 2 tunnel, vers B MTBFs, und e^+	15.5	84.5	64.3	74.3	10.2	1.8	13.7	3.8
Cold3	Cold 2 tunnel, vers B MTBFs, conv e^+	11.8	88.2	74.5	84.5	3.7	3.3	8.5	4.0
Cold4	Cold 1 tunnel, vers B MTBFs, und e^+	25.1	74.9	54.3	64.3	10.6	1.4	23.7	10.0
Cold5	Cold 1 tunnel vers C MTBFs, und e^+	15.1	84.9	64.1	74.1	10.8	1.2	13.9	3.9

Table 4.4.1.2: Version A MTBFs that were needed to get the Warm 2 tunnel LC with a undulator e^+ source to a budgeted downtime of 15%. Note that this level of detail has only been done for the linac and main damping rings. The cost increase for the first 8 klystrons and related hardware assumes the RF power sources are made redundant with high power switches to select which source is used. All other cost increases use the crude power law described in the text.

Device	Nominal MTBF (hours)	Factor improvement over nominal needed	Source of nominal MTBF	# of devices	Rough total cost for devices (% of total LC)	Very rough cost increase (% of total LC)	Increase in down time (%) if MTBF is 10 x worse
all water cooled magnets	1×10^6	10	SLAC SLC had 5×10^5 Fermilab main injector had 2×10^6	2500	2.1	1.0	3.0
Large power supply controllers	1×10^5	10	SLAC SLC had 8×10^4	1900	.39	.19	2.6
Large power supplies	2×10^5	10	Fermilab main injector had 6×10^4. TESLA design with redundant regulators estimated at 2×10^5	1900	.72	.35	3.4
Linac controls local backbone	1×10^5	15	Commonly used number for an electronics module	2200	.23	.14	1.8
Flow switches	2.5×10^5	5	SLAC SLC had 2.2×10^5	4800	.1	.04	0.9
AC power distribution - small	3.6×10^5	3	SLAC SLC had 3.6×10^5	830	.02	.004	2.1
DR kicker pulsers	7×10^3	3	SLAC SLC had 7×10^3	4	.12	.027	0.8
first 5 klystrons and related hardware (should be done with redundancy)	varied	5	varied	16		.02	2.1
Total cost						1.8	

Table 4.4.1.3: Version B MTBFs that were needed to get the Cold 2 tunnel LC with an undulator e^+ source to a budgeted downtime of 15%. Note that this level of detail has only been done for the linac and main damping rings. Note that for device = "energy overhead", the energy overhead was increased rather than make all the devices contributing to the energy more reliable. Also, for the sensitivity check in the last column, the energy overhead was decreased by 2%. The cost of increasing that energy overhead is directly estimated without the use of the power law formula. The cavity tuners and cavity piezo tuners designs both require opening the cryostat to effect repairs and had over 50 failures per year. This is an unreasonable amount of work even for the 3 month shutdown. The tuners will either have to be made very reliable (probably via redundancy) or their failure prone components made replaceable without warm-up. . The cost increase for the first 5 klystrons and related hardware assumes the RF power sources are made redundant with high power switches to select which source is used.

Device	Nominal MTBF (hours)	Factor improvement over nominal needed	Source of nominal MTBF	# of devices	Rough total cost for devices (% of total LC)	Very rough cost increase (% of total LC)	Increase in down time (%) if MTBF is 10 x worse
all water cooled magnets	1×10^6	10	SLAC SLC had 5×10^5 Fermilab main injector had 2×10^6	2800	1.4	.69	5.0
Large power supply controllers	1×10^5	40	SLAC SLC had 8×10^4	600	.10	.08	1.2
Large power supplies	2×10^5	10	Fermilab main injector had 6×10^4. TESLA design with redundant regulators estimated at 2×10^5	600	.13	.06	1.9
All electronics modules	1×10^5	3	Commonly used number for electronics modules	25000	.81	.18	3.8
Linac controls local backbone	1×10^5	9	Commonly used number for electronics modules	600	.05	.02	0.8
Vacuum valve controllers	1.9×10^5	5	SLAC SLC had 1.9×10^5 for valves + controllers. Most failures were the controllers	300	.05	.02	1.3

Continued from previous page

Device	Nominal MTBF (hours)	Factor improvement over nominal needed	Source of nominal MTBF	# of devices	Rough total cost for devices (% of total LC)	Very rough cost increase (% of total LC)	Increase in down time (%) if MTBF is 10 x worse
Flow switches	2.5×10^5	10	SLAC SLC had 2.2×10^5	1700	.03	.02	1.8
Water instrumentation	3×10^4	3	SLAC SLC had 3.5×10^4 Fermilab main injector had 5.6×10^4	330	.02	.003	1.2
AC power distribution - small	3.6×10^5	10	SLAC SLC had 3.6×10^5	700	.02	.008	1.1
first 5 klystrons and related hardware (should be done with redundancy)	varied	20		10		.08	0.4
Cavity tuner - see caption for details	1×10^6	50	SLAC SLC magnet movers had 5×10^5. Assume tuner is similar as it is a mechanical stepping motor	18000	.63	.63	0.1
Vacuum pumps on the insulating vacuum	1×10^5	6	guess	150	.03	.01	1.6
Linac energy overhead	2%	1%	Energy overhead increased from 2% to 3%			.32	5.0
Total cost						2.1	

Table 4.4.1.4: Version C MTBFs that were needed to get the Cold 1 tunnel LC with a undulator e^+ source to a budgeted downtime of 15%. Note that this level of detail has only been done for the linac and main damping rings. Note that for device = "energy overhead", the energy overhead was increased rather than make all the devices contributing to the energy more reliable. Also, for the sensitivity check in the last column, the energy overhead was decreased by 2%. The cost of increasing that energy overhead is directly estimated without the use of the power law formula. The cavity tuners and cavity piezo tuners designs both require opening the cryostat to effect repairs and had over 50 failures per year. This is an unreasonable amount of work even for the 3 month shutdown. The tuners will either have to be made very reliable (probably via redundancy) or their failure prone components made replaceable without warm-up. The cost increase for the first 5 klystrons and related hardware assumes the RF power sources are made redundant with high power switches to select which source is used.

Device	Nominal MTBF (hours)	Factor improvement over nominal needed	Source of nominal MTBF	# of devices	Rough total cost for devices (% of total LC)	Very rough cost increase (% of total LC)	Increase in down time (%) if MTBF is 10 x worse
all water cooled magnets	1×10^6	40	SLAC SLC had 5×10^5 Fermilab main injector had 2×10^6	2800	1.4	1.3	1.7
Large power supply controllers	1×10^5	80	SLAC SLC had 8×10^4	600	.1	.11	0.5
Large power supplies	2×10^5	10	Fermilab main injector had 6×10^4. TESLA design with redundant regulators estimated at 2×10^5	600	.13	.06	2.1
DR corrector power supplies	4×10^5	3	SLAC corrector and quad supplies had 4.3×10^5.	1300	.1	.02	1.1
Vacuum pump power supplies for pumps on RF power couplers	1×10^5	3	Commonly used number for electronics modules	1500	.29	.06	0.2
All electronics modules	1×10^5	3	Commonly used number for electronics modules	25000	.81	.18	5.1
Linac controls local backbone	1×10^5	30	Commonly used number for electronics modules	600	.05	.03	0.3
Linac controls sector backbone	1×10^5	9	Commonly used number for electronics modules	100	.02	.008	0.6
Linac controls timing backbone	1×10^5	9	Commonly used number for electronics modules	100	.02	.008	0.6
DR controls local backbone	1×10^5	6	Commonly used number for electronics modules	200	.03	.01	1.1

Continued from previous page

Device	Nominal MTBF (hours)	Factor improvement over nominal needed	Source of nominal MTBF	# of devices	Rough total cost for devices (% of total LC)	Very rough cost increase (% of total LC)	Increase in down time (%) if MTBF is 10 x worse
Vacuum valve controllers	1.9×10^5	30	SLAC SLC had 1.9×10^5 for valves + controllers. Most failures were the controllers	300	.05	.03	0.5
Flow switches	2.5×10^5	10	SLAC SLC had 2.2×10^5	1700	.03	.02	1.9
Water instrumentation	3×10^4	3	SLAC SLC had 3.5×10^4 Fermilab main injector had 5.6×10^4	330	.02	.003	2.5
AC power distribution - small	3.6×10^5	10	SLAC SLC had 3.6×10^5	700	.02	.008	2
first 5 klystrons and related hardware (should be done with redundancy)	varied	20		10	.08	.05	1
All klystrons and related hardware in DR (should be done by adding extra klystrons and cavities)	varied	10		6	.71	.36	0.7
Kicker pulsers	7×10^3	10	SLAC SLC had 7×10^3	84	.06	.03	3.2
Power coupler interlock sensors	1×10^6	5	Guess of 1×10^5. The increased a factor of 10 from assumed redundant design	18000	.29	.1	0.5
Power coupler interlock electronics	1×10^6	5	1×10^5 is commonly used number for electronics module. Assume redundant design improves that a factor of 10	18000	.29	.1	0.5
Vacuum valve	1×10^6	3	SLAC SLC had 1.9×10^5 for valves + controllers. Most failures were the controllers	300	.08	.02	1.9
Water pumps	1.2×10^5	3	SLAC SLC had 1.2×10^5	330	.26	.05	1.2

Continued from previous page

Device	Nominal MTBF (hours)	Factor improvement over nominal needed	Source of nominal MTBF	# of devices	Rough total cost for devices (% of total LC)	Very rough cost increase (% of total LC)	Increase in down time (%) if MTBF is 10 x worse
Vacuum mechanical devices	1×10^5	5	SLAC SLC had 7×10^4	20	.02	.005	0.7
Cavity tuner - see caption for details	1×10^6	50	SLAC SLC magnet movers had 5×10^5. Assume tuner is similar as it is a mechanical stepping motor	18000	.63	.63	0.1
Cavity piezo tuner - see caption for details	1×10^6	50	guess	18000	.63	.63	0.1
Vacuum pumps on the insulating vacuum	1×10^5	6	guess	150	.03	.01	0.2
Linac energy overhead	2%	6%	Energy overhead increased from 2% to 8%			1.9	0.8
Total cost						5.1	

Runs Warm2 and Cold2 start from the conditions of runs Warm1 and Cold1 and increase some MTBFs in order to achieve the goal 15% downtime. This could be done by increasing them all by a factor of 4. Rather than this, devices which contributed more than most to the downtime were improved by larger factors, leaving the MTBFs of most of the devices at their nominal values. It is presumably cheaper to redesign or make redundant a small fraction of the devices than to do that for all of them. The goal is to get the desired downtime for as little additional expense as possible. As the design of the LC progresses, it will be necessary to do real engineering to estimate the cost of the improvements to MTBFs and modify the goal values arrived at here in order to minimize the cost. For simplicity here, all devices which contributed more than about 0.2% to 0.3% to the downtime had their MTBFs improved sufficiently to get below that level. An exception was those components whose failure simply reduced the energy gain in the main linac. If too much downtime was caused by these components, then the energy overhead was increased instead of increasing the MTBFs, as that few percent increase in linac length would probably be cheaper than improving so many components.

While what has been improved in the simulation is MTBFs, one could get some of the gain by decreasing the MTTRs. This optimization will be left to a later stage of the design. It should be noted that decreasing the MTTR by a factor of two is not as effective as increasing the MTBF by a factor of two for devices in the accelerator tunnel because of the overhead for access and the possibility of overlapping multiple repairs.

While many MTBFs are left unimproved compared to present practice, engineers must remain vigilant. The LC is enormously complex and it will not take many design errors to make it too unreliable.

Tables 4.4.1.2 and 4.4.1.3 show the MTBFs that were adjusted to achieve the goal downtime of 15% for the warm and cold LCs (neglecting the additional 10% downtime reserved for contingency). To allow easy comparison to present state of the art, achieved numbers from some present accelerators are also given. While some of the MTBFs are rather large, this should not cause immediate panic. Assigning a 50 million hour MTBF does not mean the device must continue working for 571 years without being touched. It means that with proper preventive maintenance, which can include completely replacing the device before it wears out or replacing redundant components after they have failed, there will be an average of one failure per 571 device-years of running.

It is difficult to estimate the MTBF of a device accurately, even using historical data, as the answer varies widely with the person who did the analysis and with the running period that was considered. Industry tables of MTBFs show a variation of over a factor of 10 for the same type of device. For these reasons, simulations were done to study the effect on overall LC availability if the actual MTBF of a device were significantly worse than specified. The last column in the three tables above shows the effect on overall downtime if the particular device type had an MTBF a factor of ten worse than the budget allocated (the first times the second column). Typically the downtime increases a few percent, which indicates that failure to meet one or two of the MTBF budgets would not be a disaster. This provides some flexibility as in a project this size some mistakes will inevitably be made. The 10% availability contingency was intended to cover such problems.

The next-to-last column in the tables gives an incredibly crude estimate of the extra cost to reach the desired MTBF. Better estimates would require real engineering studies of the details of each system, analyzing the failure modes and devising remedies. This simply was not within the scope of this task force, but will be required as the LC heads towards construction. In some cases, the extra cost is only in the engineering design, which is amortized over a large number of components. In other cases, it may be necessary to provide redundancy which could significantly increase the device cost.

The crude estimates were based on the following logic:

- Since most of the devices to be improved are those that traditionally cause downtime in existing accelerators, many years of effort have already gone into making them reliable. Hence, there is likely no cheap solution to get a large improvement (like adding an extra cooling fan to a power supply).

- When a factor of 10 improvement is needed, it is assumed to be accomplished by making the most troublesome half of the device redundant, leaving about 10% remaining single point failures. In this model, a factor of 10 improvement increases the cost by 50%.
- A factor of 2 improvement might be gained by better quality control, burn-in, shake testing etc. This factor of 2 might add 13% to the cost.
- These two cost increases can be described by a power law. Each extra factor of 10 multiplies the cost by 1.5. Conveniently, the cost increase chosen to give a factor of 2 improvement lies on this power curve so we have
- The factor by which the cost increases = (MTBFdesired / MTBFnominal)$^{(\ln(1.5)/\ln(10))}$

This is obviously not universally correct, and is more likely to apply to power supplies and interlocks than to magnets and vacuum valves.

Run Warm3 (Cold3) is intended to show the difference between using an undulator e^+ source and a conventional e^+ source. Otherwise, it is identical to Warm2 (Cold2). Since the positron source was not modeled in detail and the same overall MTBF was used for both systems, the difference is not due to the specific components of the two sources. The significant change is that the undulator source requires high energy electrons before positrons can be produced, while the conventional source does not. This one difference changes the downtime from 15% to about 11% because of the shorter recovery time when both systems can be restored independently. More importantly, the amount of time spent integrating luminosity increases from about 75% to over 85% in the simulation, or from 65% to 75% including the 10% downtime contingency. This increase is due to the reduced downtime mentioned above, and to a decrease in the time spent in scheduled MD (1-4%) because more MD can be done opportunistically. If differences between the actual reliabilities of the sources are ignored, the choice of an undulator source in order to potentially produce polarized positrons reduces the integrated luminosity by more than 15%.

This large reduction in luminosity with the undulator positron source comes from three effects, all due to the need for high energy electrons in order to produce positrons. The first effect has nothing to do with reliability. It is the simple fact that with a conventional e^+ source one can do MD simultaneously in for example the e^- DR and the e^+ linac, which is not possible with the undulator. The second effect is that recovery from a downtime is slower as one cannot start tuning the e^+ system until beam is recovered in the e^- linac. The third effect is that with a conventional e^+ source, but not with the undulator source, one can do opportunistic MD in the e^+ system when something in the e^- system is down or recovering.

This problem might be ameliorated by using a conventional e^+ source to provide positrons to the e^+ DR when no undulator e^+ were available, but this case has not been simulated. The impact would depend on the switchover time between the two e^+ sources and on whether the conventional source was full power or only low power.

Run Cold4 examines the one tunnel variant of the cold LC with everything else the same as Cold2. With one tunnel, devices like linac klystrons and magnet power supplies are located in the accelerator tunnel and require access to fix. The modulators remain accessible and there are long pulsed cables connecting them to the klystrons. With the same MTBFs, the 15% downtime for Cold2 becomes 25%. Since the lumped components are causing about 10% of this, the linac and DR components must be improved by another factor of 3 on average, a total of a factor of 12 over nominal. Run Cold5 includes the MTBF improvements shown in Table 4.4.1.4, which were designed to meet this tougher demand. Not surprisingly, the list in Table 4.4.1.4 is much longer and the numbers are much higher than for the two tunnel cold version (Cold2) in Table 4.4.1.3. The difference in the total cost given in the two tables gives an idea of the incremental cost to achieve the same availability with one tunnel as with two. This difference is about 3% of the total project cost. This should be balanced against the savings in other costs such as tunneling to evaluate the cost effectiveness of a single tunnel. It is worth mentioning that there are other issues with a single tunnel, such as the extra difficulty involved in debugging problems that involve beam without access to the electronics.

Figures 4.4.1.1-4.4.1.3 illustrate the downtime budgets for the warm and cold reference design options, and for the cold one-tunnel variant. Roughly a third of the downtime is caused by the linac and DRs that have been modeled in detail. The other two-thirds is caused by site-wide systems (power, global controls, and cryogenic plants) and by the regions of the accelerator not modeled in detail. For the warm case, the linacs cause more downtime than the damping rings while that situation is reversed for the cold case. This is reasonable given the relative complexity of the damping rings. In all cases, the electron regions cause more downtime than the corresponding positron regions. This is not due to the intrinsic reliability of the components (which are the same for the e^+ and e^- systems). Rather, with the use of the undulator e^+ source, it takes longer to recover from an outage of an electron region than from the corresponding positron region as beam must be recovered first through the electron region and then through all the positron regions.

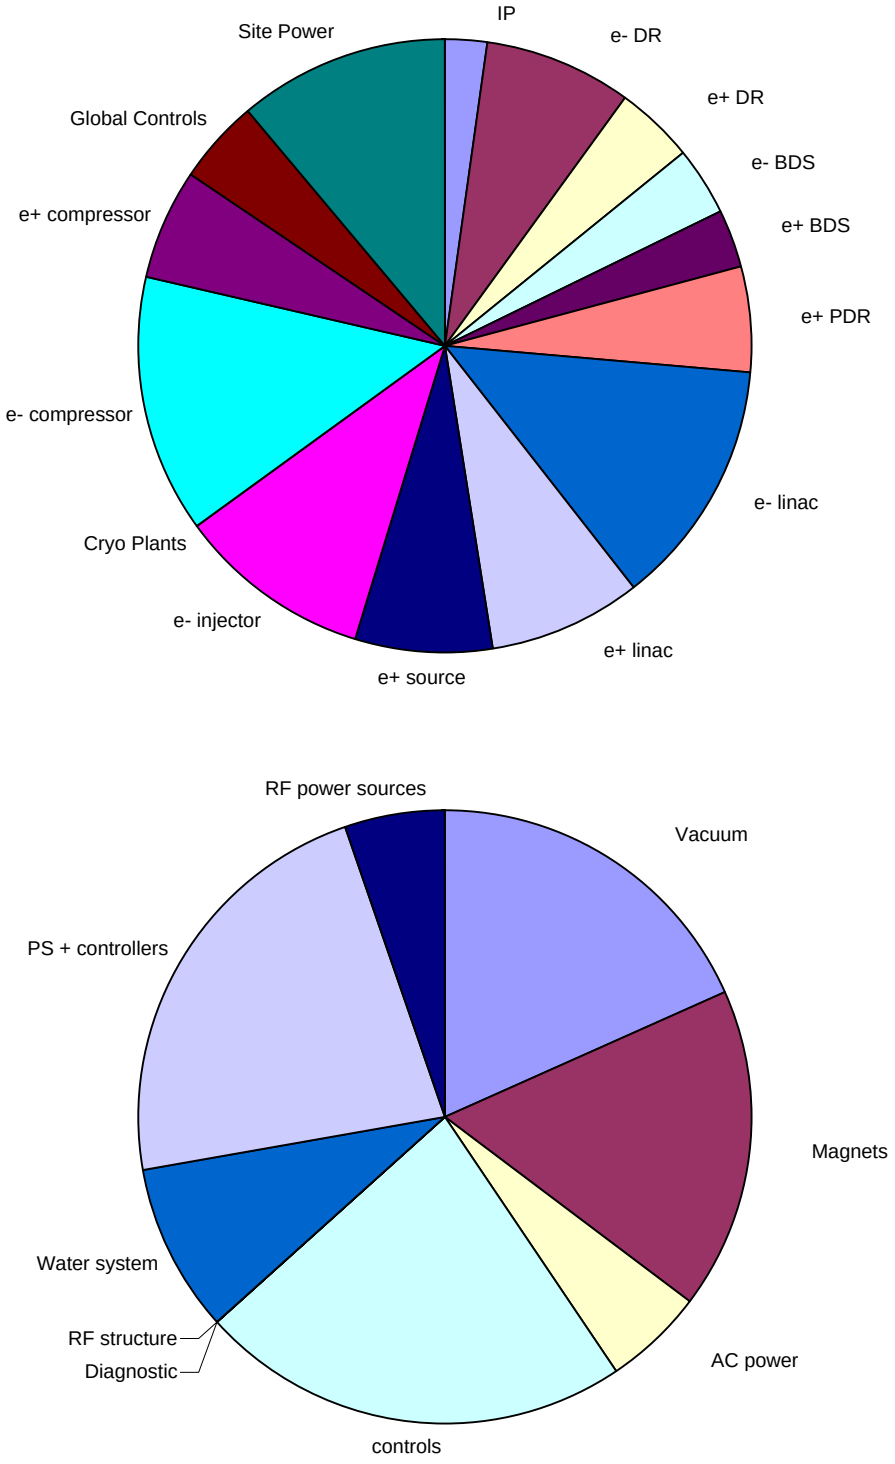

Figure 4.4.1.1: Downtime summary for run Warm2 (2 tunnel, undulator e^+ source, version A MTBFs). The top chart divides it by region while the bottom chart divides the linac and DR downtimes by system.

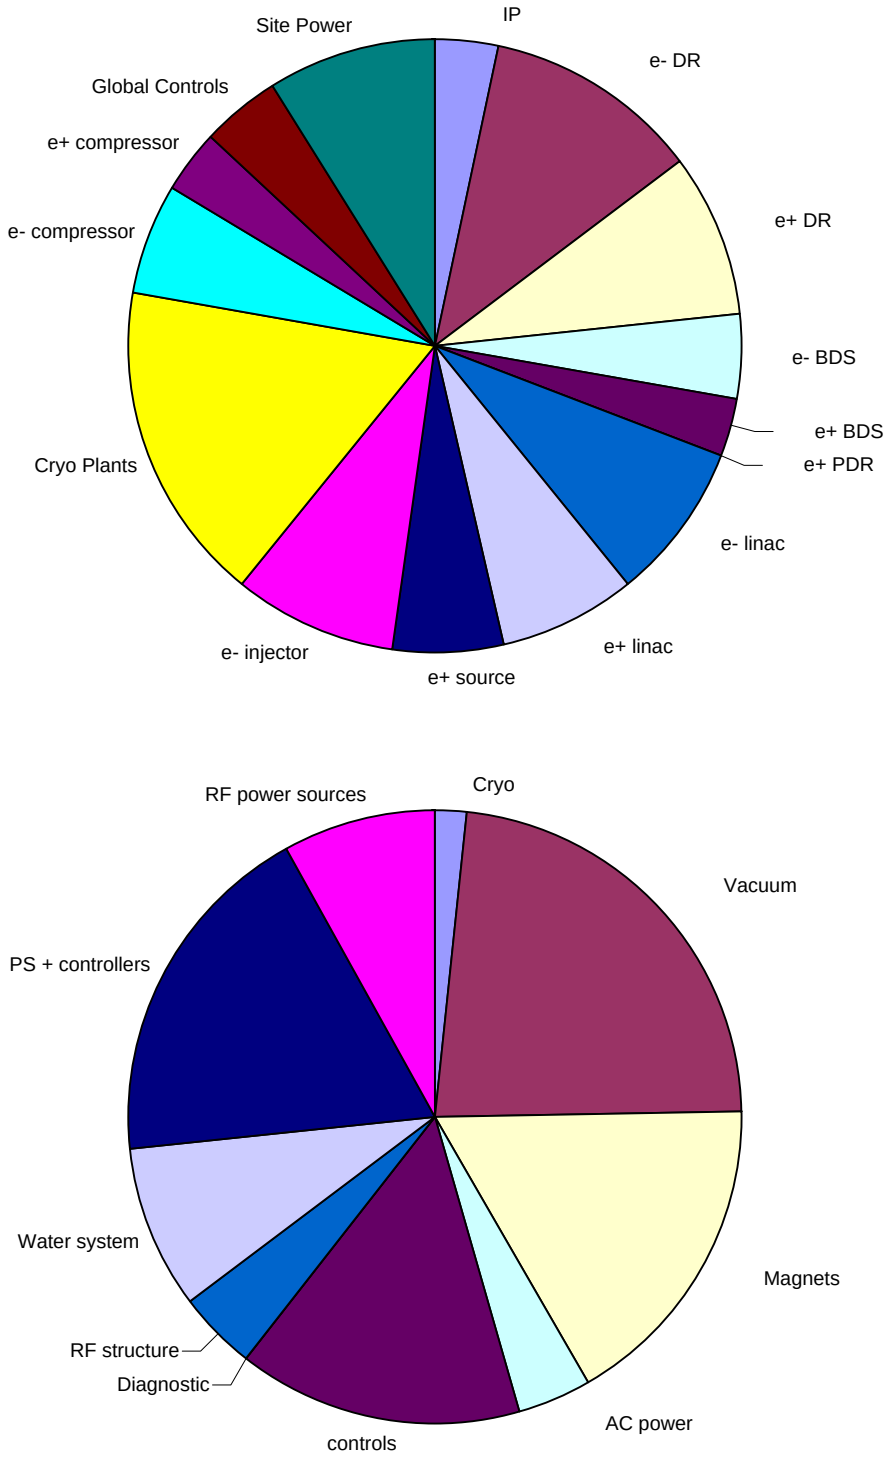

Figure 4.4.1.2: Downtime summary for run Cold2 (2 tunnel, undulator e^+ source, version B MTBFs). The top chart divides it by region while the bottom chart divides the linac and DR downtimes by system.

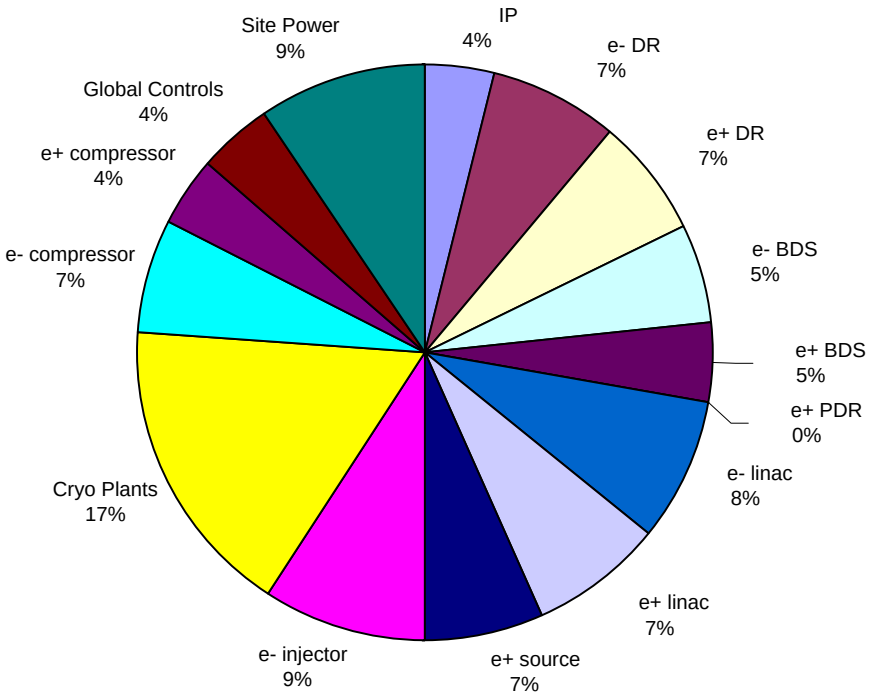

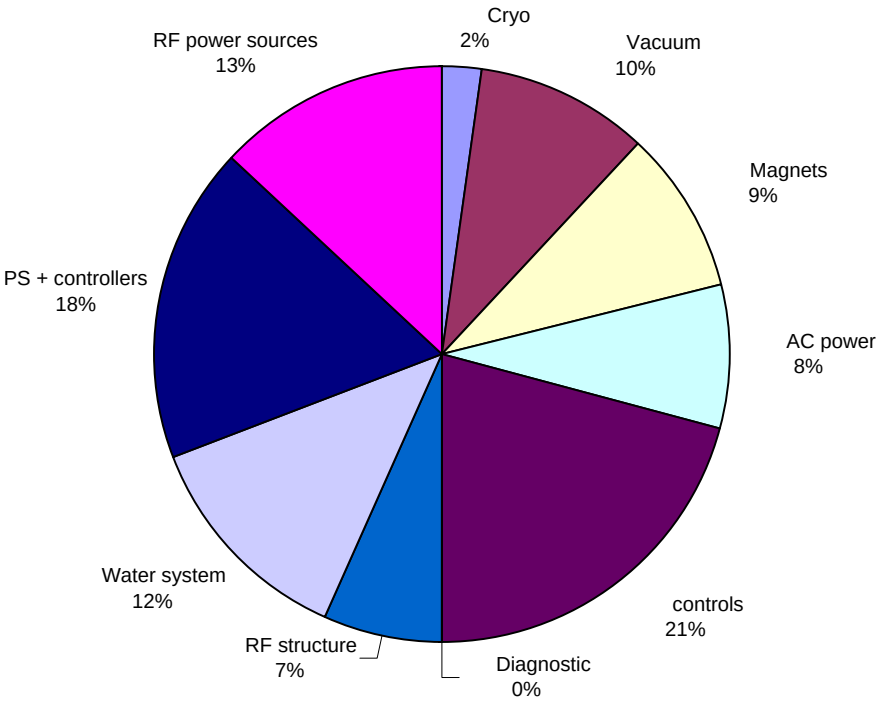

Figure 4.4.1.3: Downtime summary for run Cold5 (1 tunnel, undulator e^+ source, version C MTBFs). The top chart divides it by region while the bottom chart divides the linac and DR downtimes by system.

4.4.2 Sensitivity studies

The simulations described above have assumed conditions corresponding to steady state operation of the accelerator after most of the bugs have been ironed out and it is running fairly well. This will hopefully be the case after a couple of years of operation. To investigate what the availability might be during the commissioning period before things are running quite so smoothly, the simulations were repeated with three simple changes:

1. All the MTBFs were halved.
2. Twice as much time was spent on MD.
3. The recovery time from outages was doubled.

The results of these simulations are shown in Table 4.4.2.1. Run Warm4 (Cold6) is the same as Warm2 (Cold2) except with these "commissioning" features. It represents the base design with an undulator e^+ source. The downtime (column 3) has gone from 15% to 46-48% and the time spent integrating luminosity (column 5) has dropped precipitously from about 75% to 31%, or from 65% to 21% if the 10% downtime reserved for contingency is also included. This low an efficiency effectively means the detector gets no useful luminosity. Changing from an undulator to a conventional e^+ source makes a dramatic difference. Run Warm5 (Cold7) depicts the LC with a conventional positron source during the commissioning period. The fraction of time spent integrating luminosity improves considerably to over 64% instead of 31% (54% including contingency).

This commissioning scenario demonstrates that an LC (warm or cold) should start with a conventional positron source and only switch to an undulator source as an upgrade after the accelerator has been running for several years. Otherwise, the loss in integrated luminosity is at least a factor of two.

The next three sensitivity studies were only done for the warm design as the results were expected to be very similar for the cold design. Run Warm6 explores the sensitivity of the simulation (and presumably the accelerator) to the recovery time. It is the same as Warm2 but the recovery and tuning after a downtime take half as long. The downtime decreased from 15% to 10% which is a significant change. Achieving a short recovery time is worth considerable effort. This effort could include: controlling temperatures so they are the same during running and during access; automated setup and tuning procedures and feedbacks; careful checks and crosschecks and procedures for work that is done during an access; diagnostics to find what is broken without beam; and magnets which are covered so they can remain on during access.

Run Warm7 checks the effect of increasing the cooldown time from 1 hour to 3 hours. Otherwise, it is the same as Warm2. This change may be necessary if the radioactive air from the tunnel is going to be exhausted into a populated area (as assumed in the TESLA TDR). The downtime increased from 15% to 16% indicating that this is a small but noticeable effect, probably worth further study.

Run Warm8 checks the effect of decreasing the MTTRs by a factor of two. The downtime goes from 15% to 9% indicating that decreasing the mean time to repair is another very worthwhile effort.

Run Cold8 documents the effect of putting the cold damping ring in its own tunnel (the warm DR is always in its own tunnel). It has the same run conditions as Cold2 except that the DR is in its own PPS zone, which makes it possible to have beam in the DR while people are in the linac tunnel. Downtime decreased by only 0.5%. This is a surprisingly small effect and merits further study to make sure nothing important has been missed.

Table 4.4.2.1: Summary of sensitivity studies done to explore how simulation parameters effect the results. Each line represents one run of the simulation. The 10% downtime that is held as contingency was not simulated and hence is only reflected in the table in the column marked, incl. 10% contingency. Note that columns 3 and 4 should add to 100% as should columns 6-9. The four runs, labeled "commissioning" try to simulate the early running of the LC by halving all the MTBFs, doubling the recovery times, and doubling the amount of MD.

Run number	LC type	Simulated % time down incl. forced MD	Simulated % time fully up integrating lum or sched MD	Total % time integrating lum, incl 10% contingency	Simulated % time integrating lum	Simulated % time scheduled MD	Simulated % time actual opportunistic MD	Simulated % time useless down	Simulated # of accesses per month
Warm4	Warm2 but "commissioning"	46.2	53.8	21.3	31.3	22.4	3.6	42.7	3.2
Warm5	Warm3 but "commissioning"	30.5	69.5	57.5	67.5	2.0	14.0	16.5	4.1
Warm6	Warm2 but recovery time halved	10.1	89.9	69.9	79.9	10.0	3.0	7.1	2.8
Warm7	Warm2, but 3 hr cooldown time, not 1 hr	16.3	83.7	64.0	74.0	9.7	3.3	13.0	2.4
Warm8	Warm2 but all MTTRs halved	9.0	91.0	69.5	79.5	11.5	1.5	7.5	3.0
Cold6	Cold2 but "commissioning"	47.7	52.3	21.0	31.0	21.3	2.7	45.0	4.7
Cold7	Cold3 but "commissioning"	31.8	68.2	53.8	63.8	4.4	9.6	22.2	6.2
Cold8	Cold2, but with DRs in separate tunnels	15.0	85.0	65.4	75.4	9.6	2.4	12.6	3.8

4.4.3 Estimates of repair personnel

The simulations also explored how many people are necessary for accelerator repairs. Three classes of people were included:

1. Those who go into the accelerator enclosure to make repairs. This was limited to 50 in all the simulations.
2. Those who make repairs to devices outside the accelerator enclosure that can only be done when the accelerator is down. This was limited to 100 in all the simulations.
3. Those who make repairs to devices that are hot swappable. This was limited to 50 in all the simulations.

These numbers are for personnel required at the time of the repairs. The number of available people was chosen to be large for the simulations so that limited manpower would not degrade the availability. To see how many people are actually needed, the simulation has been run repeatedly for different numbers of staff people. This was done individually for each of the three classes of people. Figure 4.4.3.1 shows these curves for people who go into the accelerator tunnel for each of the three baseline accelerators. Table 4.4.3.1 summarizes the number of staff needed on shift by giving how many people of each class are needed to keep the downtime only 1% above its minimum value (when there are a very large number of people). Very few people of class 2 are needed as most devices outside the accelerator tunnel are hot swappable and most of the remaining devices completely break the accelerator (which means they get repaired immediately and hence are not queued up for repair when something else breaks).

The numbers from Table 4.4.3.1 cannot be directly used to determine the required staffing level at the future LC laboratory. Firstly, the simulation assumes the people are available 24/7. That requires five employees for each one on shift. Secondly, the people who repair hot swappable devices could probably be used to perform the downtime tasks. Thirdly, there are many activities that are not simulated which occur on a down day which also require people. Examples include minor upgrades, preventive maintenance, and minor repairs which don't degrade the accelerator performance but should be fixed (e.g. a small water leak).

Table 4.4.3.1: The number of each class of person that must be available 24 hours a day to avoid degrading the downtime by one percent.

Class of people	Warm 2 tunnel	Cold 2 tunnel	Cold 1 tunnel
In accelerator tunnel	4	6	25
Outside accelerator tunnel	4	4	4
Repair hot swappable devices	6	2	Not simulated

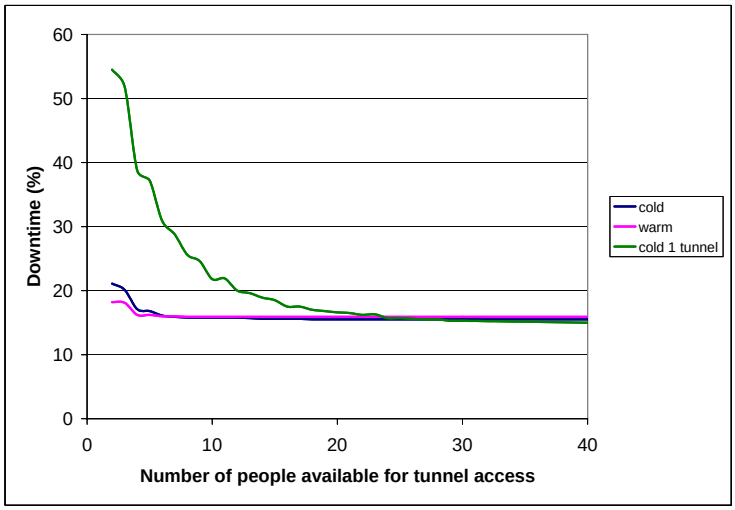

Figure 4.4.3.1: Downtime as a function of the number of maintenance people available to enter the accelerator tunnel to perform repairs.

4.5 Summary

A simulation has been developed which calculates the overall downtime of an accelerator given the reliability characteristics of its parts. This has been used to develop a detailed availability budget for the linac and damping rings and a less detailed budget for the other regions of the LC.

This budget is very much a first iteration. The allocation of the unavailability is not fully optimized. However, the budget does give an indication of the level of difficulty that will be involved to make the LC sufficiently reliable. The tools developed are very useful and should continue to be used in parallel with actual reliability engineering to refine this first iteration. It would also be useful to benchmark the results against a real accelerator.

The power of the simulation really shows up when comparing different versions of the accelerator. These comparisons shouldn't depend heavily on the details of the availability budget. The comparisons reveal:

1. The fact that an undulator positron source requires well tuned high energy electrons before positrons can be produced significantly reduces the integrated luminosity of a LC. For example in the warm LC after a few years of running, the luminosity integrated in a year would be 18% less for an undulator positron source than for a conventional one.

2. During commissioning, it would be far worse; a factor of two less luminosity would be integrated.

3. There is not a great difference between warm and cold 2 tunnel designs. Both are very large and complex accelerators where significant effort and expense will be needed to make them reliable enough.

4. If all components had the same reliability, the cold 1 tunnel design would have a downtime of 25% instead of the 15% of the 2 tunnel design. The required improvements to Mean Time To Failures (MTBFs) of many components and the necessary 5% increase in the energy overhead needed to recover the 15% availability are costly. This cost should be compared to the amount saved in building one tunnel instead of two to see if it is truly worthwhile.

5. The effects of Mean Time To Repair (MTTR) and the speed of recovery from a downtime are quite significant. It will be worth considerable design effort to keep these small. Success at this could reduce the requirements on the MTBFs

4.6 Addendum: Component descriptions

The linacs and damping rings were divided into many components. Below is a description of these components and the nominal MTBFs and MTTRs (in hours) that were used as starting values in the simulation.

Magnets - includes the magnet itself, interlocks, water connections and cables from power supply to the magnet, but not the power supply or its controller. All magnets require access to repair and most failures stop operation. Exceptions are correctors or trim windings where it is assumed one can tune around scattered failures. It is also assumed that one can match around quadrupoles in the cold main linac, which has relatively weak focusing. The luminosity degradation due to a matched around quad or corrector is 1%, and 0.1% for a trim.

General magnet - MTBF: 1×10^6(Water cooled) 1×10^7(superconducting), 1×10^7(permanent magnet)
General magnet - MTTR: 8(repair) 2(quad retune) 472(repair SC quad later)
Correctors & trims - MTBF: 1×10^7 MTTR: 0.5 (retune) 2 (later repair)

Kickers - in the cold damping rings, injection or extraction requires 20 out of 21 kickers so two failures are required before stopping operation. Kickers for tuneup dumps are considered part of diagnostics.

Kicker magnet - MTBF: 1×10^5 MTTR: 8
Kicker pulser - MTBF: 1×10^4 MTTR: 2

PS + controllers - power supplies and their controllers are generally installed in the support housing and do not require access for repair, except for the one tunnel variant of the cold machine where most are inaccessible during operation. Their impact depends on the type of magnet controlled as discussed under magnets. Since these are a well identified reliability risk, redundancy is already incorporated in the designs. Cold linac and DR magnets have 4 of 5 redundancy. Most warm linac and DR magnets have 1 of 2 redundancy. This raises the estimated MTBF to 2×10^5 from 5×10^4 experienced at the labs.

Power supply - MTBF: 2×10^5 (redundant) MTTR: 2 (normal) 4 (large)
Controller - MTBF: 1×10^5 MTTR: 1

Movers - movers are installed on the beamline and require access to repair but it is assumed one can tune around failed movers. The luminosity degradation due to a matched around mover is 0.1%. Mover controllers are 16 channels, accessible and hot swappable.

Mover - MTBF: 5×10^5 MTTR: 0.5 (retune) 2 (repair later)
Controller - MTBF: 1×10^5 MTTR: 1

Magnet flow switches - water cooling systems for electromagnets require flow switches. For the warm linac and cold DR straights, there is assumed to be a flow switch every 100 m. For the warm DR and cold DR arcs, there is a flow switch every 2 girders or approximately 10 m. These are in the tunnel and require access for repair.

Water flow switch - MTBF: 2.5×10^5 MTTR:1

Vacuum mechanical device - this covers any device in vacuum which moves into the beamline such as stoppers or profile monitors. Since they are on the beamline, access is required for repair.

Vacuum device - MTBF: 1×10^5 MTTR: 8

Vacuum pumps and PS - The beamlines have vacuum pumps about every 100 m in the warm regions and and one per 9 cryomodules in the cold linacs. (During the study, one pump per module was used in error. This has little effect on the simulation results). The pumps are in the tunnel and require access to repair; the power supplies are in the support housing (except for 1 tunnel variant). Since the systems have a certain amount of redundant pumping, the luminosity is not degraded when a pump fails.

Vacuum pump - MTBF: 1×10^7 MTTR: 4

Vacuum pump power supply - MTBF: 1×10^5 MTTR: 1

Vacuum valves and controllers - There are assumed to be vacuum valves every 200 m in all regions. This number was chosen as a reasonable guess as the published numbers for the cold linac were too small to be reasonable and for the cold DR too large. A valve failure is assumed to interrupt operation (a valve stuck out would likely not be noticed). Valves require access to repair, controllers are in the support housing (except 1 tunnel).

Vacuum valve - MTBF: 1×10^6 MTTR: 4

Vacuum valve controller - MTBF: 1.9×10^5 MTTR: 2

Beamline water systems - These systems supply water to beamline components. Each pump has associated instrumentation and a flow switch. All systems are assumed to be accessible for repair (except 1 tunnel). In the DRs and cold linacs, water pump failures turn off the machine. In the warm linacs, most of the pumps supply water to the structures and hence only reduce the energy overhead and can be replaced during operation.

Water pump - MTBF: 1.2×10^5 MTTR: 4

Water pump instrumentation - MTBF: 3×10^4 MTTR: 2

Water flow switch - MTBF: 2.5×10^5 MTTR: 1

Beam position monitors - Failure of the physical cavity or stripline monitor are extremely rare and ignored. Only readout module failures are included. These are in the support housing (except 1 tunnel) and hot swappable during operation. Luminosity degradation is assumed to be 0.1% per failed unit.

BPM - MTBF: 1×10^5 MTTR: 1

Wire scanners (laser and conventional) - Failures are assumed to be in the instrumentation or in the laser itself, all accessible during operation. Luminosity degradation for laser wires is 5% as they are critical for tuning. Loss of conventional wires does not degrade luminosity.

Wire scanner - MTBF: 1×10^5 MTTR: 2

Tuneup Kickers - Like laser wires, these are important for efficient tuning. The kickers are in the tunnel and require access, the pulsers are in the support housing. Luminosity degradation is taken to be 5%.

Kicker - MTBF: 1×10^5 MTTR: 8
Kicker pulser - MTBF: 1×10^4 MTTR: 2

RF power sources - The RF power sources for DRs and linac structures and cavities are klystrons powered by a modulator for the linacs or by a power supply for the DRs. Each klystron has a pre-amplifier, vacuum gauge/controller, vacuum pump and pump power supply. All of these are accessible for repair and hot swappable, except for the 1 tunnel cold machine where the linac klystrons and all associated devices are installed in the tunnel. While MTBFs for klystrons and modulators vary slightly, the other components are similar so they are described here only once. When a klystron vacuum pump fails, the klystron must be replaced, so the MTTR is long.

Klystron - MTBF: 3×10^4 (DRs) 2.5×10^4 (warm linac) 4×10^4 (cold linac) MTTR: 8
Modulator - MTBF: 5×10^4 (cold linac) 1×10^5 (warm linac) MTTR: 4
DR Klystron power supply - MTBF: 5×10^4 MTTR: 4
LLRF controls - MTBF: 1×10^5 MTTR:1
Klystron pre-amp - MTBF:1×10^5 MTTR: 1
Klystron vacuum gauge - MTBF: 1×10^5 MTTR: 1
Klystron vacuum pump - MTBF: 1×10^7 MTTR: 8
Klystron pump power supply - MTBF: 1×10^5 MTTR: 1

RF cavities (warm DR) - There are 5 normal conducting cavities in each DR of the warm machine, of which only 4 are required for operation. Each cavity has LLRF controls. The cavity is on the beamline and requires access to repair. The controls are accessible and hot swappable. Each DR cavity is powered by a klystron. Loss of the power source turns off the associated cavity, but does not interrupt operation unless there are two failures at once.

Cavity - MTBF: 1×10^8 MTTR: 24

RF cavities (cold DR) - There are 12 superconducting cavities in each DR of the cold machine, of which only 11 are required for operation. Each cavity has LLRF controls. Four cavities are in a single cryomodule, fed from a single power coupler which has interlock sensors, coupler interlock controls and two vacuum pumps with two power supplies, all of which are required for cavity operation. The cavity, coupler and pumps are on the beamline and require access to repair. The controls and pump supply are accessible (except for 1 tunnel). The cryo system has a cryo vacuum enclosure which may develop minor leaks or more serious problems. Each DR 4-cavity cryomodule is powered by a klystron. Loss of a power source turns off the associated 4 cavities, which interrupts operation.

Cavity - MTBF: 1×10^8 MTTR: 72
Coupler - MTBF: 1×10^7 MTTR: 16
Coupler interlock sensors - MTBF: 1×10^5 MTTR:1
Coupler interlock controls - MTBF: 1×10^5 MTTR:1
Coupler vacuum pump - MTBF: 1×10^7 MTTR: 4
Coupler pump power supply - MTBF: 1×10^5 MTTR: 1
Cryo vacuum leak - MTBF: 1×10^5 MTTR: 8
Cryo vacuum failure - MTBF: 3×10^5 MTTR: 8

Water system for DR klystrons - The DR klystron cooling water system has a water pump and instrumentation. There is one system per ring for the warm machine and one per 4-cavity cryomodule for the cold machine. There are two flow switches per pump on the cold machine and one per klystron or cavity on the warm. The pumps and instrumentation are accessible as are the flow switches. A water pump outage removes enough energy to interrupt operation, but the warm machine only loses one klystron/cavity pair when a flow switch fails.

Water pump - MTBF: 1.2×10^5 MTTR: 4

Water pump instrumentation - MTBF: 3×10^4 MTTR: 2

Water flow switch - MTBF: 2.5×10^5 MTTR: 1

RF power sources and structures (warm linac) - In the warm main linac, an RF unit consists of a pair of klystrons fed from a single modulator with a single LLRF controller. Each RF unit feeds two SLED systems, which power 16 accelerator structures. Each klystron has a pre-amplifier, flow switch, vacuum gauge/controller, vacuum pump and pump power supply. The two SLED systems together have about 100 vacuum pumps driven by 5 20-channel pump controller chassis. Cooling water is supplied by one water pump with instrumentation per 16 RF units. All parts are accessible and hot swappable, except the structures.

SLEDs - MTBF: 1×10^5 MTTR: 4

SLED vacuum pump - MTBF: 1×10^7 MTTR: 4

SLED vacuum pump chassis - MTBF: 1×10^5 MTTR: 1

Structure - MTBF: 1×10^8 MTTR: 168 (1 week)

RF power sources (cold linac) - In the cold main linac, an RF unit consists of a modulator powering a klystron. Each RF unit feeds 2.5 cryomodules, each containing 12 9-cell cavities. Each klystron has a pre-amplifier, flow switch, vacuum gauge/controller, vacuum pump and pump power supply. Cooling water is supplied by one water pump with instrumentation per 2 RF units. In the two tunnel layout, all parts are accessible and hot swappable, except the cryomodules. In the single tunnel layout, only the modulator is accessible. Everything else is located in the accelerator housing and requires access for repair. In addition, the klystron in the single tunnel is connected to the modulator by high power pulsed cables that are up to 2.5 km long, with a pulse transformer located next to the klystron. The MTBF is taken from experience with electrical power distribution cables (unpulsed but similar voltage). While there are spare cables, they must be routed to the correct RF unit so the MTTR is also long.

Pulsed cables (1 tunnel) - MTBF: 2×10^5 MTTR: 8

Pulse transformer (1 tunnel) - MTBF: 2×10^5 MTTR: 4

RF cavities (cold linac) - In the cold main linac, there are 12 9-cell cavities per cryomodule, and a total of 30 cavities per klystron/modulator RF unit. Since the cavities, tuners and couplers are in cryomodules, repairs require warmup of the cryo string and hence are only feasible during the yearly shutdown. Rather than turn off the entire RF unit when one cavity has a problem, there are a number of possible softer failure modes. Because of the complexity of the possible responses, many of these failure modes have been modelled individually.

Cavity gradient - If the cavity trips at full gradient, the response is to lower the gradient for just that cavity up to 10 MV/m by a combination of detuning and adjustments to the 3 stub tuner. If the cavity cannot stay on at lower gradient, the RF unit is turned off until the next access when the individual cavity can be disconnected.

Cavity gradient degraded (average loss 8 MeV) - MTBF: 1×10^8 MTTR: 672 (4 wks)
Cavity broken/disconnected (loss 29 MeV) - MTBF: 1×10^8 MTTR: 672 (4 wks)

Cavity tuner and drive - The tuner is required to keep the cavity on resonance and compensate for pressure and thermal changes. It is in the cryomodule so repair is long. Failures might not affect performance for a while but eventually the cavity would have to be disconnected to prevent errant behavior.

Cavity tuner and drive (loss 29 MeV) - MTBF: 1×10^6 MTTR: 672 (4 wks)

Piezo tuner - The piezo tuner is required to compensate for Lorentz force detuning at gradients above 24 MV/m. For redundancy, two tuners per cavity are planned. Failures require lowering gradient to 24 MV/m. At a nominal gradient of 28 MV/m, this can be done for a single cavity. For operation at 35 MV/m, either all cavities would need to operate at somewhat lower gradient or the coupler for the failed cavity would have to be disconnected.

Cavity piezo tuner (loss 5 MeV) - MTBF: 5×10^6 MTTR: 672 (4 wks)

LLRF - Complex LLRF algorithms are required to protect the cavity. Failure means that cavity must be detuned.

LLRF (loss 29 MeV) - MTBF: 1×10^5 MTTR: 1

Coupler breakdown - If the coupler shows breakdown activity, the gradient must be lowered. A temporary fix without access is to lower the gradient for the entire RF unit. If the problem is not too serious, lowering to 20 MV/m should be adequate. For serious failures, the entire RF unit must be turned off. During the next access, the single cavity can be disconnected and full power RF restored to remaining cavities. The coupler is in the cryomodule so repairs require a long shutdown.

Coupler problem (minor) (average loss 240 MeV) - MTBF: 1×10^7 MTTR: 2 (disconnect)
Coupler problem (major) (loss 872 MeV) - MTBF: 1×10^7 MTTR: 2 (disconnect)
Coupler repair - MTTR: 672 (4 wks)

Coupler interlocks - Both the input and HOM couplers are instrumented with a variety of sensors and interlocks. Some sensors will be redundant and some cross checks with other sensors may be used to backup a failed sensor. If there is no workaround for the failed sensors, the klystron must be turned off. The sensors are on the cryomodule and require access but not warm-up for repair, the electronic readout for the sensors is in the support tunnel, except for 1 tunnel variant. Long MTBFs are used to account for the redundancy that is assumed.

Coupler interlock sensors (loss 872 MeV) - MTBF: 1×10^6 MTTR: 1
Coupler interlock electronics (loss 872 MeV) - MTBF: 1×10^6 MTTR: 1

Coupler vacuum pump - There is a vacuum manifold for each cryomodule to distribute vacuum to all the couplers. Only a single pump is required but two pumps and power supplies are installed for redundancy. Failures require the entire RF unit to be turned off. The pumps are in the accelerator housing and require access to repair. The pump power supplies are in the support tunnel (except with 1 tunnel). The redundant pump/supply MTBFs are assumed to be 10 times nominal.

Coupler vacuum pump (loss 872 MeV) - MTBF: 1×10^8 MTTR: 4

Coupler pump power supply (loss 872 MeV) - MTBF: 1×10^6 MTTR: 1

Linac cryomodules (cold linac) - In the cold main linac, each cryo plant feeds two cryo units (each about 2.5 km long). Insulating vacuum is divided into 5 sections per cryo unit (each about 500 m long). These sections are separated by vacuum barriers . Small leaks to the insulating vacuum may be overcome by hooking up a local turbo pump during an access. Major problems require warmup. 10 cryomodules are grouped into a cryostring with valves and controls at the ends of the string, in particular a JT valve. Repairs to any of these systems require access. Failures require that the entire cryo system affected be turned off. The JT valve has a relatively short MTTR as it is assumed there will be some workaround short of replacement.

Cryo vacuum enclosure (no luminosity) - MTBF: 3×10^5 MTTR: 8

Cryo insulating vacuum (loss 3488 MeV) - MTBF: 1×10^5 MTTR: 8

Cryo JT valve (loss 3488 MeV) - MTBF: 3×10^5 MTTR: 2 (workaround)

Electrical distribution - A simple model was used to develop a rough count of breakers in the electrical distribution system. High power breakers (> 0.5 MW) were assigned for the RF units. Lower power breakers (50 KW - 0.5 MW) were assigned for magnets, pumps, and controls. In the DRs, there is a high power circuit to drive each klystron and a low power circuit per klystron for controls. In the warm linac, there is a high power circuit per four klystrons and one per two klystrons in the cold linac. There are also an equal number of control circuits. Low power control circuits were assigned for every 10 magnets, for each structure water pump in the warm linac, and for each controls sector.

High power electrical circuit - MTBF: 3.6×10^5 MTTR: 4

Low power electrical circuit - MTBF: 3.6×10^5 MTTR: 2

Controls, MPS and PPS - A simple model was used to develop a rough count of units for the control system, Machine Protection and Personnel Protection. First, failures of these systems which affect the entire complex were assigned an overall downtime of 0.2%. This downtime was treated explicitly like site power or cryogenics, rather than assigned to a particular region. The controls system was assumed to have a communications backbone every 200 m, designated a "sector". In the linacs, each sector also has a timing distribution associated with it. Each sector is then subdivided into local backbone units (logically crates). These are assigned 1 per klystron in the linacs, and 1 per 10 magnets in the DRs. Individual controllers for magnets, movers and LLRF have been included earlier. In addition, each klystron is assumed to have a local timing distribution and 3 more controls modules associated with it. For Personnel Protection, there is assumed to be a circuit for each entry point, 5 in the linacs, 1 for the warm DR, 2 for the 2 arcs of the cold DR. Global communications within a region such as Machine Protection, Fast Feedforward on beam current in the linacs, and fast feedback in the DRs were counted as two network circuits per region. Failures at the backbone level or higher interrupt operation in that region. Klystron control failures only take the affected units offline. All of these controls are accessible for repair without access, except for the single tunnel linac where the sector and local controls are in the accelerator tunnel.

Sitewide controls, MPS, PPS - MTBF: 2.5×10^3 MTTR: 5

Sector controls backbone - MTBF: 1×10^6 MTTR: 2

Sector timing distribution - MTBF: 1×10^5 MTTR: 1

Local controls backbone - MTBF: 1×10^5 MTTR: 1

Local timing distribution - MTBF: 1×10^5 MTTR: 1

Local klystron controls - MTBF: 1×10^5 MTTR: 1
PPS controls - MTBF: 1×10^5 MTTR: 1
MPS, Fast Feedforward, Fast Feedback networks - MTBF: 5×10^3 MTTR:1

Bibliography

[TDR] R. Brinkmann et. al., eds., "TESLA Technical Design Report", Part II, DESY-2001-011, (March, 2001)

[ZDR] T.O. Raubenheimer, ed. [NLC ZDR Design Group], "Zeroth order design report for the Next Linear Collider," SLAC-R-0474, (1996)

[NLC01] N. Phinney, ed. [NLC Collaboration], "2001 report on the Next Linear Collider: A report submitted to Snowmass '01," SLAC-R-571, (2001)

[NLC03] NLC 2003 Configuration

[TRC] G. Loew et. al., "International Linear Collider Technical Review Committee: Second Report", SLAC-R-606, (2003)

[JLC] N. Toge, ed. [JLC Design Study Group], "JLC design study," KEK-REPORT-97-1 (1997).

[ISG] C. Adolphsen, et. al., [International Study Group Collaboration], "International study group progress report on linear collider development," SLAC-R-559, (2000)

[GLC] K. Abe et. al, "GLC Project Report", KEK Report 2003-7, Sep. 2003. Also available at http://lcdev.kek.jp/ProjReport/

[1] DOE/NSF High Energy Physics Advisory Sub-panel on Long Range Planning for U.S. High Energy Physics, (Jan., 2002)

[2] "Design Considerations for an International Linear Collider", the American Linear Collider Physics Group Executive Committee, (2003)

[3] "Parameters for the Linear Collider", Parameters Subcommittee, International Linear Collider Steering Committee, Sept. 30, 2003,

[4] Daniel Schulte, et. al., "Tests of Three Linear Collider Beam Dynamics Simulation Programs", LCC-0091 (2002).

[5] T.O. Raubenheimer, P. Tenenbaum, "Brief Review of Linear Collider Beam-Based Alignment", LCC-0129 (2003).

[6] N. Walker, et. al., "Damping Ring to Interaction Point Beam Transport Issues," p.657, *Proceedings of the 2003 Particle Accelerator Conference* Portland, OR (2003).

[7] P. Tenenbaum, R. Brinkmann, V. Tsakanov, "Beam-Based Alignment of the TESLA Main Linac," pg. 515, *Proceedings of the 2002 European Particle Accelerator Conference*, Paris (2002).

[8] D. Schulte, N. Walker, "Simulations of Static Tuning for the TESLA Collider", pg. 2736, *Proceedings of the 2003 Particle Accelerator Conference*, Portland, OR (2003).

[9] Paul Emma, [ZDR, Chapter 5] and subsequent undocumented studies in 1998.

[10] C. Sramek, et al., "Beam-Beam Interaction Simulations with Guinea Pig", LCC-0125 (2003).

[11] Y. Nosochkov, et al., "Tuning knobs for the NLC final focus", pg. 476, *Proceedings of the 2002 European Particle Accelerator Conference*, Paris (2002)

[12] Mark Woodley, Andrzej Wolski, "The NLC Main Damping Ring Lattice," LCC-0113, (February, 2003)

[13] Ina Reichel, Andrzej Wolski, "A Reduced Emittance Lattice for the NLC Positron Predamping Ring," LCC-0114, (February, 2003)

[14] Y. J. Batygin, J. C. Sheppard, "Start-to-End Simulation of Polarized Positron Collection in the NLC", LCC-0117, SLAC (2003).

[15] A. Seryi, talk presented at NLC MAC meeting, (June, 2003).

[16] B. Parker, talk presented at NLC MAC meeting, (November, 2002).

[17] J. Volk, et al., "Adjustable Permanent Quadrupoles for the Next Linear Collider," pg. 217, *Proceedings of the 2001 Particle Accelerator Conference*, Chicago, IL (2001).

[18] Y. K. Batygin and J. C. Sheppard, "Capture and polarization of positrons in a proposed NLC polarized positron source," SLAC-PUB-9891 and pg. 2751, *Proceedings of the 2003 Particle Accelerator Conference*, Portland, OR (2003).

[19] Y. Nosochkov, private communication.

[20] Y. Batygin, J. C. Sheppard, "Thermal Stress Analyses for a Thermal Multislug Beam NLC Positron Target," LCC-0090, SLAC (2002).

[21] R. Rimmer, et al., "An RF Cavity for the NLC Damping Rings", pg. 924, *Proceedings of the 2001 Particle Accelerator Conference*, Chicago, IL (2001).

[22] R.A. Rimmer, D. Atkinson, J.N. Corlett, G. Koehler, D. Li, N. Hartman, J. Rasson, T. Saleh, R. Weidenbach,"RF Cavity R&D at LBNL for the NLC Damping Rings, FY2000/2001," LCC0072, (Oct., 2001).

[23] S. De Santis, "Coupled Bunch Instabilities in the NLC Damping Rings," LCC0069, (May, 2001).

[24] Tor Raubenheimer, Andrzej Wolski, "Comparison of Alignment Tolerances in the Linear Collider Damping Rings with those in Operating Rings," LCC-0112, (January, 2003).

[25] A. Wolski, private communication.

[26] P. Emma, T. Raubenheimer, "Circumference Correction Chicanes for Damping Rings," LCC-0008, (June, 1998).

[27] Paul Emma, "Cost and Performance Optimization of the NLC Bunch Compressor Systems," LCC-0021, (August, 1999).

[28] Paul Emma, "More Options for the NLC Bunch Compressors," LCC-0035, (January, 2000).

[29] S. Marks, private communication.

[30] A. Drozhdin, L. Keller, T. Markiewicz, T. Marayuma, M. Mokhov, T. Raubenheimer, A. Seryi, P. Tenebaum, M. Woodley, "NLC Collimation Update Study," LCC-0118, (May, 2003).

[31] P. Raimondi and A. Seryi, "A Novel Final Focus Design for Future Linear Colliders," LCC-0039, (May, 2000).

[32] P. Raimondi and A. Seryi,"Novel Final Focus Design for Future Linear Collider", Phys. Rev. Lett. **86**, 3779 (2001).

[33] S. Smith, "Design of an NLC Intra-pulse Feedback System", LCC-Note-56, SLAC, March 2001.

[34] P. Burrows, "Feedback on Nanosecond Timescales (FONT): Results from First Beam Tests at the NLCTA at SLAC", pg. 687, *Proceedings of the 2001 Particle Accelerator Conference*, Chicago, IL (2001)

[35] J. C. Sheppard, "Conventional Positron Target for a Tesla Formatted Beam", LCC-0133, (SLAC-TN-03-072), (November, 2003).

[36] The radiation spectrum from planar and helical undulators is described in NLC Notes LCC-085 and LCC-095.

[37] T. A. Vsevolozhskaya, A. A. Mikhailichenko, E. A. Perevedentsev, G. I. Silvestrov and A. D. Chernyakin, "Helical Undulator For Conversion System Of The VLEPP Project," SLAC-TRANS-0225, presented at Int. Conf. on High Energy Accelerators, Novosibirsk, USSR, Aug 7-11, 1986; A. Mikhailichenko, "Short-period SC undulator," CBN-03-4 (2003).

[38] A. Mikhailichenko, "Pulsed undulator for test at SLAC of the polarized positron production," CBN-03-5 (2003).

[39] W. Stein, J. C. Sheppard, "Structural Modeling of the TESLA TDR Positron Target", LCC-0089, SLAC (2002)

[40] W. Stein, J. C. Sheppard, "NLC Polarized Positron Photon Beam Target Thermal Structural Modeling", LCC-0087, SLAC (2002)

[41] V.K. Bharadwaj, et al., "Analysis of Beam-Induced Damage to the SLC Positron Production Target", pg. 2123, *Proceedings of the 2001 Particle Accelerator Conference*, Chicago, IL (2001).

[42] LCC-note in preparation on Ti damage, (2003).

[43] J. Sheppard, private communication, (2002).

[44] Y. J. Batygin, J. C. Sheppard, "Start-to-End Simulation of Polarized Positron Collection in the NLC", LCC-0117, SLAC (2003).

[45] J.P. Labrie, H. Euteneuer, "Power Handling Capability of Water-cooled CW Linac Structures", Nucl. Instrum. Meth. **A247**, 281 (1986).

[46] K. Flottmann, V. Paramonov, eds., "Conceptual design of a Positron Pre-Accelerator for the TESLA Linear Collider, "DESY TESLA-99-14, (1999).

[47] K. Flottmann, private communication, (2003).

[48] D. P. Pritzkau and R. H. Siemann, "Experimental study of RF pulsed heating on oxygen free electronic copper," Phys. Rev. ST Accel. Beams **5**, 112002 (2002).

[49] For an example of a flexible spin rotation system, see [ZDR, 5.3.3].

[50] For an example of an energy compressor, see [ZDR, 4.3.3].

[51] C. Santelli et. al., "Technical Layout of the TESLA Damping Ring", LNF-01/003(NT) (2001).

[52] B.I. Grishanov, F.V. Podgorny, J. Ruemmler, V.D. Shiltsev, "Very Fast Kicker for Accelerator Applications", DESY TESLA-96-11, (1996).

[53] E. Chojnacki and J. Sears, "Superconducting RF Cavities and Cryogenics for CESR III Upgrade", Cornell SRF 990716-09, (1999).

[54] T. Tajima, et al., "The Superconducting Cavity System for KEKB", pg. 440, *Proceedings of the 1999 Particle Accelerator Conference*, New York, NY (1999).

[55] W. Decking, G. Hoffstatter, T. Limberg, "Bunch Compressor for the TESLA Linear Collider", DESY TESLA-00-40, (2000).

[56] M. D. Woodley, P. E. Emma,"Measurement and Correction of Cross Plane Coupling in Transport Lines", pg. 196, *Proc. 20th Int. Linac Conference*, Monterey, CA, Aug. 2000; SLAC-PUB 8581, (2000).

[57] A. Drozhdin, G. Blair, W. Kozanecki, L. Keller, T. Markiewicz, T. Marayuma, M. Mokhov, O. Napoly, T. Raubenheimer, A. Seryi, D. Schulte, P. Tenenbaum, N. Walker, M. Woodley, F. Zimmerman,"Comparison of the TESLA, NLC and CLIC beam collimation system performance", LCC-0111, SLAC (January, 2003).

[58] S. Claudet, Ph. Gayet, Ph. Lebrun, L. Tavian and U. Wagner, "Economics Of Large Helium Cryogenic Systems: Experience From Recent Projects At CERN", LHC Project Report 317, (December 1999).

[59] S. Wolff, H. Lierl, B. Petersen, H. Quack, "The Tesla Cryogenic Distribution System", TESLA Report 2001 - 37.

[60] S. Claudet, Ph. Lebrun, L. Tavian,"Towards Cost-To-Performance Optimization of Large Superfluid Helium Refrigeration Systems ", LHC Project Report 391, (2000).

[61] T. R. Strobridge and D. B. Chelton, "Size and Power Requirements of 4.2 K Refrigerators", Advances in Cryogenic Engineering, 12, (1967).

[62] T. R. Strobridge, "Refrigeration for Superconducting and Cryogenic Systems", IEEE Trans. Nuclear Science, Vol. NS-16, (1969).

[63] T. R. Strobridge, "Cryogenic Refrigerators–An Updated Survey", NBS Technical Note 655, (1974).

[64] M. A. Green, R. A. Byrns and S. J. St. Lorant, "Estimating the Cost of Superconducting Magnets and the Refrigerators Needed to Keep Them Cold", Advances in Cryogenic Engineering 37, pg. 637, (1991).

[65] R. A. Byrns and M. A. Green, "An Update On Estimating the Cost of Cryogenic Refrigeration", Advances in Cryogenic Engineering 43, p 1661, (1998).

[66] H. Edwards, P. Garbincius, et al., "Report on the TESLA Engineering Study/Review", Fermilab TM-2179, (2002).

[67] DOE Budget Formulation Handbook, pg. 11.

[68] NLC 1999 DOE Project Review, SLAC, (May, 1999).

[69] "Advanced Photon Source Conceptual Design Report", ANL-87-15, April, 1987.

[70] "ITER Final Design Report", http://www.naka.jaeri.go.jp/ITER/FDR/, July 2001.

[71] Seryi, A., LCD Group Meeting - "Update on NLC Stability Studies", NLC Group Contribution to TRC, October 31, 2002.

[72] Frame, P., 2000, "Geologic Conditions at Representative Sites for the NLC in California," Prepared for: Stanford Linear Accelerator Center, 30p.

[73] Lettis & Associates, Inc., Unruh, J., Lettis, W.R., 2002, "Phase II Fault and Seismic Hazards Investigation"–North of Delta Offstream Storage Integrated Storage Investigations, California Department of Water Resources Storage Investigation, August 2002.

[74] Robert A. Bauer and David L. Gross, "Geology of the Greater Fermilab Region", Illinois State Geological Survey.

[75] S. Smith, "Cavity Beam Position Monitor for the Next Linear Collider", LC-02, SLAC, (2002).

[76] C. Montag, private communication.

[77] V. Shiltsev, private communication.

[78] A. Wolski, W. Decking, "Comparison of Emittance Tuning Simulations in the NLC and TESLA Damping Rings," TESLA 2002-10, DESY (2002).

[79] M. Boege, B. Keil, T. Schilcher V. Schlott, "Commissioning of the Fast Orbit Feedback at SLS", pg. 3386, *Proceedings of the 2003 Particle Accelerator Conference*, Portland, OR (2003).

[80] J. Sekutowicz et. al, "Cold and Beam Tests of the first prototypes of the superstructure for the TESLA Collider", pg. 467, *Proceedings of the 2003 Particle Accelerator Conference*, Portland, OR (2003); J. Sekutowicz et. al, "Test of two Nb superstructure prototypes", Phys. Rev. ST Accel. Beams **7**, 012002 (2004).

[81] N. Walker, D. Schulte, G. White, "Intra-Bunch-Train Luminosity Optimisation for the TESLA Linear Collider", pg. 2733, *Proceedings of the 2003 Particle Accelerator Conference*, Portland, OR (2003).