Availability Task Force Report 'Appendix 1'

Availability Task Force, December 7, 2009

1. Introduction and Goal

This appendix is the preliminary report of the availability task force. This task force was formed to find a reasonable way to make sure that SB2009 met the availability requirements of the ILC. Those requirements are unchanged from those used in the RDR:

- There should be 9 months of scheduled running and 3 months of scheduled downtime for maintenance and upgrades.
- Total unscheduled downtime should be less than 25% of scheduled runtime.
- In our simulation we allow only 15% unscheduled downtime. The other 10% is held as contingency for things that were missed in the simulation or for major design or QA problems.

So far the task force has concentrated on the use of a single tunnel for the linac using either the DRFS or Klystron Cluster RF systems. This is one of the biggest cost savings in SB2009 and the biggest availability concern. That study is relatively complete. Other changes included in SB2009, (the central region integration and positron source changes), have had a more cursory examination and look OK, but more work is needed.

We took a three pronged approach to the problem.

- We refined the design of the RF systems to provide the goal availability. This included making some components redundant, ensuring the DRFS klystrons could be replaced rapidly during a scheduled maintenance day and changing the run schedule from a 3 month long downtime to two one-month downtimes with a 24 hour preventive maintenance period scheduled every two weeks during the run.
- We gathered and updated information on reliability of components (mean time between failures)
- We integrated all this information into AVAILSIM, the program which simulates breakdowns and repairs of the ILC to learn how much downtime there was and what the causes were. We then adjusted some inputs (length of preventive maintenance period, energy overhead and longer MTBFs) to achieve the required availability.

The availability task force has only studied the direct effects of the SB2009 changes on the availability. Other effects of the change to a single tunnel such as fire safety; need for space to install extra equipment in the accelerator tunnel; cost; installation logistics; radiation shielding of electronics and the

effect of residual single event upsets; and debugging of subtle electronics problems without simultaneous access to the electronics and beam have been or will need to be studied by other groups.

 Note that the important result of our work is not the availability we have achieved but rather the design that was necessary to do so. The simulation acts as a guide that allows us to assess design and component performance changes. The availability estimate obtained by iteration of the inputs through this process was actually pre-determined as a goal before we started work. That said, Section 2 gives the results of the simulations and summarizes the conclusions, Section 3 describes the basic machine design we modeled with emphasis on availability, Section 4 describes many of the ingredients that were needed to achieve the desired availability, Section 5 gives some details about how the simulation was done, and section 6 describes some outstanding issues.

2. Results and Summary

Our results are preliminary for several reasons. Due to time constraints we have not updated the inputs to AVAILSIM to the full SB2009 design. While the linac modeling is fairly accurate, we still have 6 km DRs and the injectors are not in the BDS tunnel, (see Central Region Integration section), for most of the simulation runs. There are also several details we are not pleased with and hope to improve (such as the cryoplants and AC power disruptions being the largest causes of downtime or the long recovery times needed after a scheduled preventive maintenance day). Nevertheless, we believe the model is accurate enough for the differences in availability of the various machine configurations we studied to be meaningful.

Our main result is shown in Figure 1 which for four different RF schemes shows the simulated unscheduled downtime as a function of the main linac energy overhead. This is a very important figure, so it is worth making very clear what is meant by it.
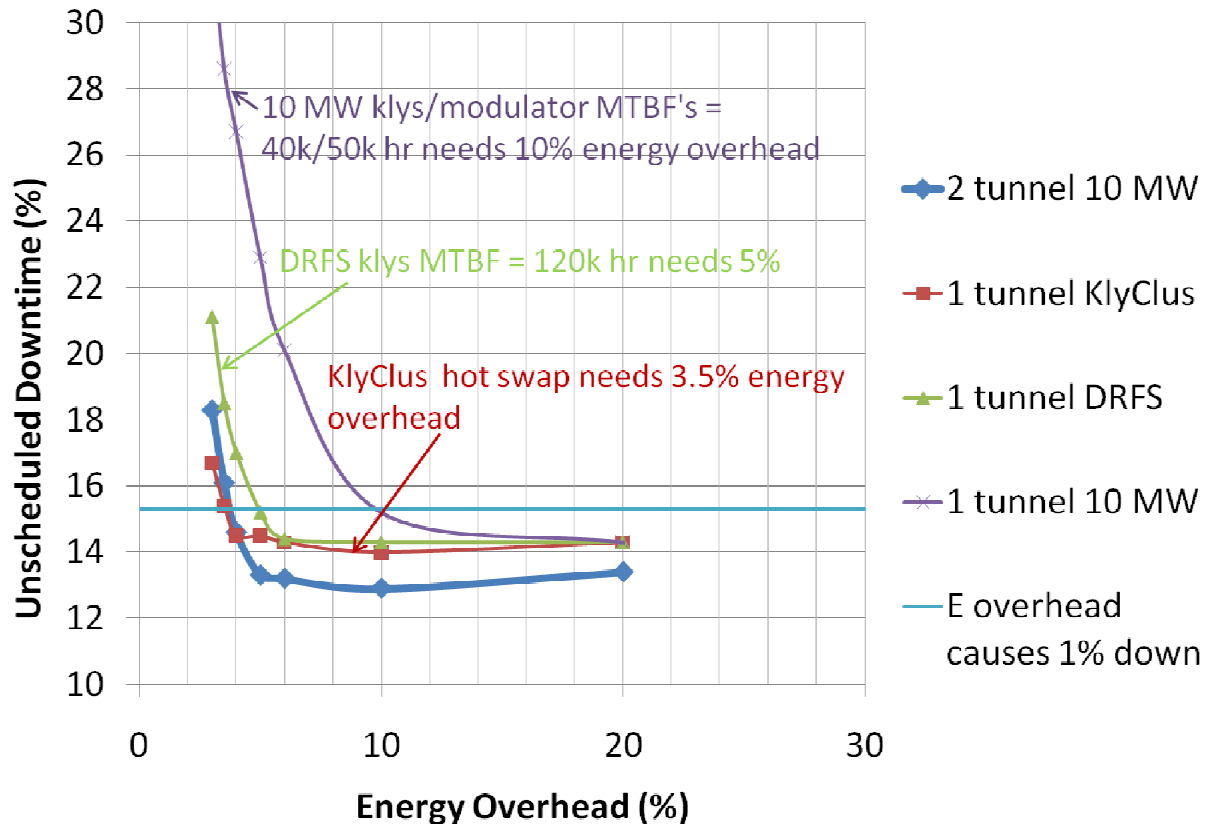
Figure 1. The unscheduled downtime as a function of the main linac energy overhead

- An example of the meaning of the horizontal axis is that for 250 GeV design energy and 10% energy overhead, the beam energy, if everything were working perfectly, would be 275 GeV. The plot is mainly relevant when we are trying to run at the full design energy. If for physics reasons we are running at lower energies, then the effective energy overhead is much greater and downtime due to inadequately working RF systems would be much less than shown. The steep rise at the left side of the plot comes about because a small number of RF sources, klystron or modulator, fail quite quickly after starting up. Repair takes time, either because entry to the linac tunnel is required (distributed RF system) or because it simply takes time to do the work (klystron cluster / reference design).

- Four different accelerator configurations were simulated resulting in the 4 lines shown. The general trend for each line is that as energy overhead is decreased, the unscheduled downtime increases. This is expected as less acceleration related items need to fail before the design beam energy cannot be reached and hence the accelerator must be shut down for unscheduled repairs.

- Note that the three one linac tunnel designs all have the same unscheduled downtime when there is 20% energy overhead. This is such a large energy overhead that failures of components whose failure reduce the energy cause no downtime. The residual 14.5% downtime is caused by things like magnet power supplies, controls and AC power disruptions.
- For illustrative purposes, a horizontal line is drawn 1% above the points at 20% energy overhead. This represents what the total unscheduled downtime would be if 1% downtime were caused by broken components (e.g. klystrons or cavity tuners) making it so the design beam energy cannot be reached. One can then see how much energy overhead each design needs to avoid causing this much downtime.
- At high energy overhead, the 1 tunnel designs have about 1% more downtime than the 2 tunnel design. This represents a doubling of the downtime caused by linac components and a 50% increase for the RTML components that share the linac tunnel. This increase is expected as a typical mean time to repair for a component in a support tunnel is 1-2 hours. When that component is in the accelerator tunnel (as is the case for the 1 tunnel designs) an extra hour is added to allow radiation to cool down and a second hour is added to allow the accelerator to be secured, turned back on and magnets standardized. Note that in studies done for the RDR, the total downtime doubled when we went from two tunnels to one. What saves us from that fate in this study is that only the linac was changed from two tunnels to one. In all other areas devices like power supplies are modeled as accessible when the beam is on.
- The one tunnel 10 MW design degrades fastest with decreasing energy overhead probably due to the 40k and 50k hr MTBFs assumed for the klystron and modulator.
- DRFS does better probably due to the redundant modulator and 120k hour klystron MTBF assumed.
- KlyClus does still better due to the ability to repair klystrons and modulators while running and the internal redundancy built into a cluster. (A single klystron or modulator can die and the cluster can still output its design RF power.)

The conclusion that can be drawn from Figure 1 is that either approach (Klystron Cluster or the Distributed RF System) would give adequate availability with the present assumptions. The Distributed RF System requires about 1.5 percent more energy overhead than the Klystron Cluster Scheme to give the same availability for all other assumptions the same. This small effect may well be compensated by other non availability related issues. With the component failure rates and operating models assumed today, the unscheduled lost time integrating luminosity with a single main linac tunnel is only 1% more than the two-tunnel RDR design given reasonable energy overheads. Note that all non-linac areas were modeled with support equipment accessible with beam on.

Figure 2 and Figure 3 show details about what area of the accelerator and what technical systems caused the downtime in a typical simulation run (KlyClus with 4% energy overhead). Note that the

unscheduled downtime is distributed over many areas and technical systems. No one thing dominates. In particular, the linacs and long transport lines of the RTML cause 20% of the total downtime of 15% or a total of 0.20*0.15 = 3% of the scheduled running time. Note that the cryogenic system plants and site power are the technical systems that cause the most downtime. Both of these are modeled as single components in the simulation with availabilities based on aggressive estimates from experts familiar with performance of such systems at existing labs. They are both worth further study.
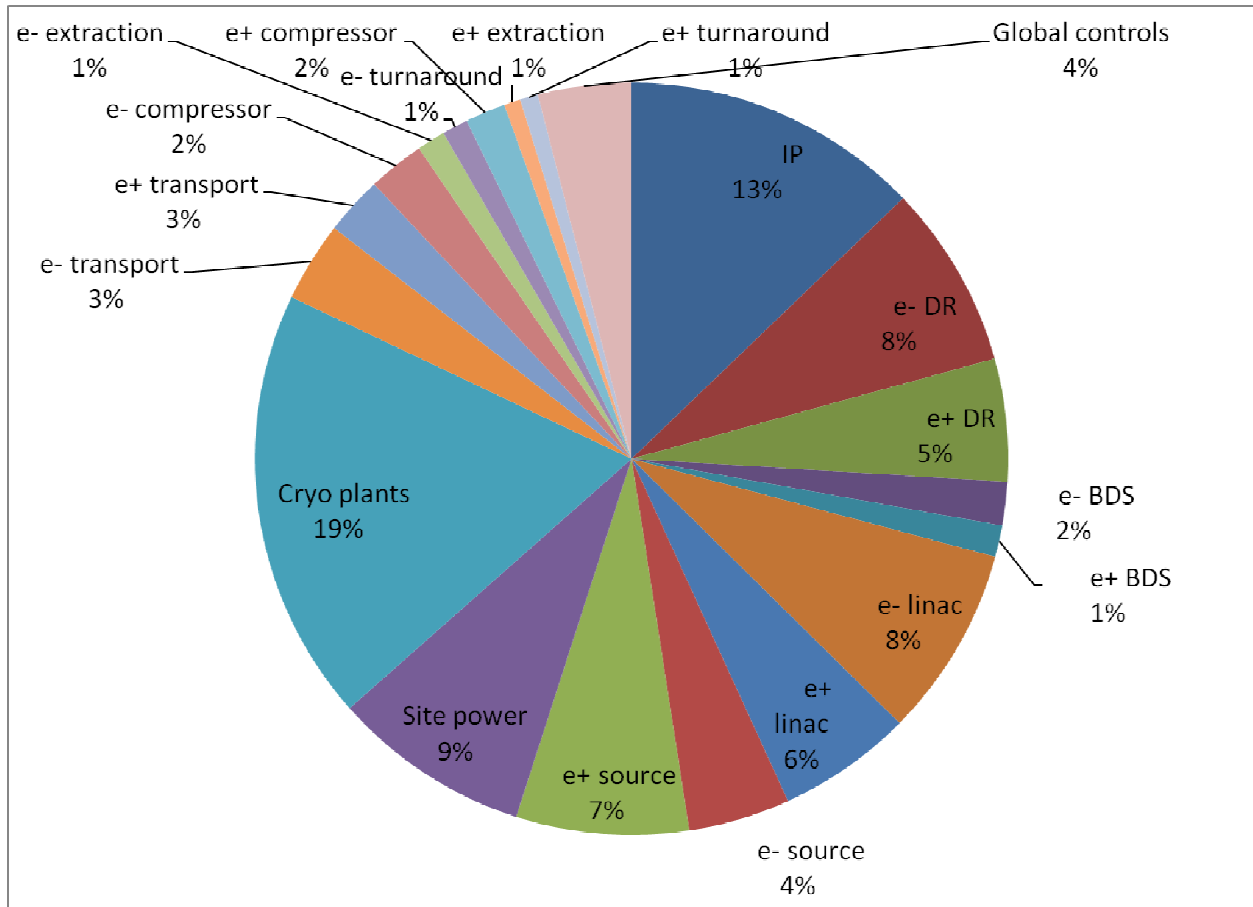

Figure 2 shows the distribution of the downtime by area of the accelerator for a typical simulation run (KlyClus with 4% energy overhead). The downtime fractions shown are percent of the total downtime of about 15%. So the actual downtime caused by the cryo plants is 19% of 15% = 2.8%. The IP does not really cause 13% downtime. That is actually the excess time spent recovering from scheduled maintenance days.
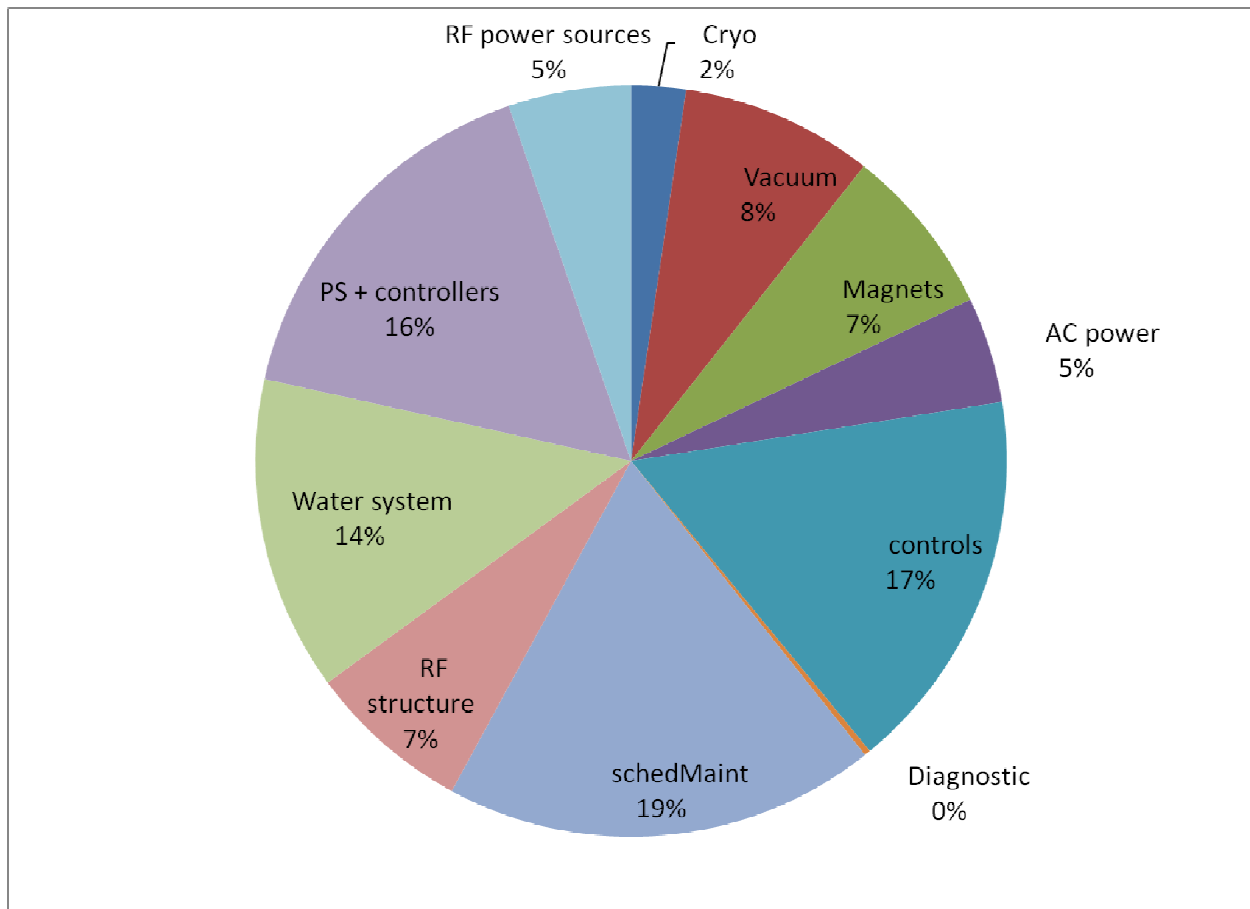
Figure 3 shows the distribution of the downtime by technical area for a typical simulation run (KlyClus with 4% energy overhead). The downtime fractions shown are percent of the total downtime of about 15%.

As explained in the introduction, our goal was to find design ingredients that achieved the desired availability goal. This section shows we have achieved that goal. The next two sections describe the ingredients used to attain that goal. You will see it requires different construction and operations methods than are typically employed for HEP accelerators. Methods used by light sources are needed. Component mean time between failures must be comparable to the best achieved at any accelerator and some are factors of more than 10 better than those achieved at major HEP labs. Preventive maintenance is crucial as is built-in redundancy.

Preliminary conclusions of impact of single main linac tunnel on availability and Recommendation:

- **The assumptions made to obtain the desired availabilities for all designs are quite aggressive and considerable attention will have to be paid to availability issues during design, construction and operation of the ILC to achieve the simulated availabilities.**

- **The RF power system as described in the RDR is unsuitable for a single linac tunnel design as there is a significant decrease in availability without further improvements in MTBF's, an increase in energy overhead and/or changes in maintenance schedules.**

- There are two alternate RF power system designs proposed for single tunnel linac operation. (The Klystron Cluster and the Distributed RF System). Either approach would give adequate availability with the present assumptions. The Distributed RF System requires about 1.5 percent more energy overhead than the Klystron Cluster Scheme to give the same availability for all other assumptions the same. This small effect may well be compensated by other non availability related issues.

- With the component failure rates and operating models assumed today, the unscheduled lost time integrating luminosity with a single main linac tunnel is only 1% more than the two tunnel RDR design given reasonable energy overheads. Note that all non-linac areas were modeled with support equipment accessible with beam on.

- The recommended RF overhead (see figure 1: 'unscheduled downtime as a function of the main linac energy overhead') is 3.5% . This corresponds to the nominal availability of about 15% for the Klystron Cluster HLRF scheme and slightly poorer availability for the Distributed RF HLRF scheme. **The cost impact of this recommendation is summarized in section xx.**

3. Background - changes applied to the Reference Design

In this section we describe the machine model used to evaluate the availability with a focus on differences with respect to the 2007 Reference Design.

3.1. Machine Model

Improvements proposed to the Reference Design are expected to affect the estimated accelerator availability. In order to estimate availability in the new design, both the basic machine model and the assumptions that support the simulation, (for example, operations and maintenance models), need review. Through this process, the availability of the complex was also re-examined and re-optimized.

The changes foreseen for the new baseline are described in the Overview, Section xx. The part of the new baseline with the greatest impact on availability is the removal of the support tunnel for the main linac and RTML. A single tunnel variant, in which no support tunnel space was allocated to any ILC subsystem, was studied during the development of the Reference Design. The estimated availability of that variant was shown to be poor unless there were very substantial improvements in component performance. The configuration studied for the proposal described here allows un-obstructed access to source, DR and BDS support equipment. Only the main linac and RTML support tunnel is removed.

3.2.Operations and Maintenance Model

The ILC is scheduled to operate 9 months of the calendar year, with the remaining 3 months allocated for machine shutdown and preventative maintenance. In the RDR the assumption was for a single 3—month shutdown. We have explored other operating models where the 3 months of total scheduled downtime is distributed throughout the calendar year. It should be apparent that less RF power overhead would be needed, (depending on how long it takes for replacement), if there is a shorter time between maintenance periods when failed klystrons are replaced. For this report, we have chosen a model with two 1-month shutdowns per year and one 24-hr preventative maintenance period every two weeks. All failed klystrons in the main linac will be replaced during the 24-hr maintenance days, so 100% of the installed RF power is available immediately following a maintenance day.

Availability estimates made for the Reference Design were based on experience at labs with high energy colliders. Experience with component performance and operations and maintenance practice was included in the analysis. For further development of the ILC design, we intend to take advantage of the knowledge gained in recent years at light sources and 'factories', such as the B-Factories. In general, experience at these installations has been much better than older, more established facilities. An important part of that improvement is due to changes in managing component lifetime performance and maintenance. These typically involve a more rigid industry-like quality assurance process, with a stronger reliance on scheduled downtime and less emphasis on 'opportunistic maintenance', a practice which includes a substantial ad-hoc, rapid response maintenance and repair effort. Additionally, the use of industrial standard-practice system designs and components (section) that have good reliability performance has been assumed in the analysis.

3.3.HLRF - Klystron Cluster System (KCS) availability design

The impact of the KCS HLRF system on machine availability is considered negligible. Redundancy will be built in to the low level RF, and the high power systems contain spare sources so most failures do not cause downtime. The proposed Klystron Cluster System allows continuous access to all active RF source equipment. Further, the system includes provisions for maintenance of the high power equipment with little or no impact on the operation of the linac.

The distribution waveguide is assumed to be robust against breakdown. Any fault (e.g. breakdown or vacuum leak) in it is a single point of failure and will cause downtime. The availability simulation does not include these faults.

3.4.HLRF – Distributed RF System (DRFS) availability design

There are two main availability issues for DRFS: Radiation damage and replacement of failed HLRF components.

A large number of klystrons will be needed since one klystron feeds its power to four cavities in DRFS scheme compared to 26 cavities in RDR. Maintainability of klystrons is a key issue for DRFS. A longer MTBF for klystrons is required for DRFS availability and a redundant scheme will be adopted for PS.

One of the main availability issues for DRFS is the radiation damage of components. Almost all DRFS components will be installed in the tunnel where we need to examine the potential radiation hazard which is caused by acceleration of the dark current. The data concerning this issue is unfortunately scarce. At this moment, we follow the considerations given in case of the EuroXFEL project, and assume to introduce radiation shield of 10cm-thick concrete and 1cm-thick lead.

Another important availability issue is the maintenance of HLRF components, in particular the klystrons. A large number of klystrons are use in the main linac tunnels, as one klystron feeds its power to four cavities, as opposed to 26 cavities by one klystron in RDR. A longer MTBF for klystrons is expected to be realized with the measures as described before. However, maintainability of klystrons and other components remains to be a key issue. In the following, the maintenance scenario for the DRFS and the associated workload are examined.

As discussed in above, the operation schedule is assumed to consist of a repetition of 2-week continuous operation (312 hrs) interreupted by one-day maintenance in between. The numbers of failed components, which require replacement or repair, after each of the 2-week operation, are estimated and are summarized in Table 4.6.3.2. The MTBF assumed in the estimates are also indicated.

Table 4.6.3.2: Estimated number of HLRF components across the entire ILC requiring replacement or repair after completion of each of the 2-week operation period, assuming the listed MTBF and component count.

Component	# of units requiring replacement or repair	MTBF assumed	Total # of units deployed at the ILC
DC power supply	2	50,000 hours	325
MA modulator	1.5	70,000 hours	325
MA klystron	12	110,000 hours	4225

It should be noted that the scheme taken with redundant DC PS and MA modulator will provide substantially longer "system" MTBFs for the DC power supplies and MA modulators than what are quoted in Table 4.6.3.2. The MTBF quoted for the klystrons is based on the operational experience of KEKB linac, where 4.7 million hours of cumulative operation time has been logged at 60 klystron sockets. BI (barium-impregnated) cathode and lower cathode loading, which assure longer life time of klystron, are adopted for this project.

In order to estimate the human resources that are required during each of the one-day maintenance, the time that it takes to replace each of the components in question has been examined. Table 4.6.3.3 shows the summary of this study.

Table 4.6.3.3: Estimated times for the repair work of DRFS.

Action	Time for unit piece of work	Rationale
Transportation of klystron	0.5 person-hours / tube	2 persons in 2 hours could bring 8 tubes on one carrier.
Removal of a failed klystron and installation of a replacement klystron	4 person-hours / tube	2 hours with 2 persons
Time for personnel to move from one point of repair to another	2/3 person-hours / tube	20 minutes with 2 persons
Replacment of a MA modulator	6.67 person-hours / modulator	
Replacement of a DC power supply	27 person-hours / DC power supply	

As a consequence, the human resource that it takes to replace the 12 MA klystrons on each of the one-day maintenance day would amount to 62 person-hours. Likewise, 1.5 MA modulators will require 10 person-hours, and 2 DC power supplies will take 54 persone-hours to repair.

In summary, the required human resource for HLRF during maintenance is calculated to be 126 person-hour, which corresponds to 16 person-days. This is likely to be manageable in the light of other work that would be necessary during maintenance days.

4. Key ingredients and their impact - Estimating MTBFs

4.1.Process for developing MTBFs for the final models

The simulation program, 'Availsim' receives as input a set of component MTBFs that were derived from operational and in-the-field experience. Since one goal of the simulations was to determine a set of MTBFs that met the availability criteria for the machine design, these MTBFs were treated as the 'Starting MTBFs. The initial Availsim simulations using these MTBFs gave an estimate of the overall machine availability given these starting MTBFs, the overall machine design, and the various other input assumptions that have been described previously in this appendix.

The results of the simulations were evaluated against the machine availability criteria. A second set of MTBFs was generated through 'educated-guesswork' based on the gap between the machine availability criteria and the relative distribution of downtimes across the different technical systems and accelerator areas. A second round of Availsim simulations were performed using the updated set of MTBFs. Additional iterations were performed until the machine availability criteria were met. The Improvement Factors coming from the Availsim simulations (the ratios between final MTBFs and starting MTBFs) have triggered availability-focused R&D programs for several technical systems, including HLRF modulators, magnet systems, and Control Systems. As noted above in the Introduction and Goal section, this is the real output of the process.

It should be noted, however, that there is no unique set of 'final' MTBFs, since during the optimization process the relative MTBFs can be adjusted to apportion total downtime differently across the various technical systems.

4.2. Source of the Starting MTBFs

As noted earlier, the MTBFs comprising the Starting MTBFs for the Availsim simulations were as much as possible derived from in-the-field experience.

The set of Starting MTBFs use for the RDR simulations was largely drawn from the years of operating experience at the SLC at SLAC and the Tevatron accelerator complex at Fermilab. As such, the set of MTBFs was representative of the availability that had already been achieved at operating high energy physics accelerator facilities.

For SB2009 simulations, the Availability Working Group proposed reviewing and revising 'up' the Starting MTBFs for individual components so as to use best-in-class achieved MTBFs unless there would be a clear and significant cost penalty over comparable systems. In particular, the intention was to take into account recent experiences with the Light Sources and the B-Factories, both of which had achieved better overall availability than had been achieved at both SLC and the Tevatron. Similarly, it was desired to take into account industry and commercial MTBF data for Commercial Off the Shelf (COTS) components. By and large, each of the components fits into one of the following five categories.

- Utility or industrial equipment where there is published MTBF data based on operating experience of a very large installed base. For example, we have used the IEEE 'Gold Book' (IEEE Std 493™2007) for electrical distribution MTBFs.
- Commercial off-the-shelf commodity equipment such as computing systems and networks where there are high availability solutions routinely deployed.
- Commercial low-complexity components such as flow switches. In these cases, we have tried to de-emphasize the downtime contributions by arbitrarily assigned large MTBFs.
- Specialized or custom-designed equipment that is accelerator specific and where there is already operating experience. Our MTBF estimates are based on experience from accelerator laboratories.

Most difficult to estimate are MTBFs for new specialized equipment where as yet there is little or no operating experience. In these cases, we can only apply good engineering judgment and where appropriate make comparisons with other existing equipment.

Largely due to time constraints, there has been only limited update of the Starting MTBFs in the RDR, and to date only ten of the Starting MTBFs have been revised, leaving others as potential candidates for revision.

4.3.MTBF Data from the RDR and SB2009 Simulations

Table A1 summarizes the MTBF data from both the RDR and the SB2009 simulations. Names of the individual components are listed in Column A. Columns B-D and F-H show Starting MTBFs, Final MTBFs, and Improvement Factors for each component for the RDR and SB2009 simulations respectively. Columns J-N list several sources of MTBF field data. In Column F (Starting MBTFs for SB2009), highlighted in green are those Starting MTBFs that have been revised for the SB2009 simulations. Conversely, highlighted in gold in Column H are the Final MTBFs that have become more demanding in SB2009. It is perhaps surprising that that Final MTBFs for several components have been increased for SB2009 over for RDR, even though they are not obviously part of the configuration changes in SB2009. As noted earlier, this is not a definitive set of necessary Final MTBFs since there are different ways to apportion the total downtime across the technical systems. It should also be reiterated that the simulation results as a whole are strongly dependent on other important assumptions that are part of the Availsim model.

An indication of the degree of difficulty deemed necessary to meet the SB2009 MTBFs can be gleamed by comparing the Final MBTFs with the Input MTBFs in Columns J-N.

Table 1 - Starting and Final MTBFs for RDR and SB2009 Availsim

Device	RDR starting MTBF	RDR table A factor	RDR final MTBF	New starting MTBF	SLC MTBF	FNAL Tevatron MTBF	FNAL Main Injector MTBF	APS MTBF	other MTBF
mttf_electronic_module	1.0E+05	3	3.0E+05	1.0E+05	9.9E+03				
mttf_PS_controller	1.0E+05	10	1.0E+06	**1.1E+06**	8.0E+04	1.8E+05	1.1E+05	1.1E+06	
mttf_controls_local_backbone	1.0E+05	3	3.0E+05	1.0E+05					
mttf_magnet	1.0E+06	20	2.0E+07	**2.0E+06**	5.0E+05		2.0E+06		
mttf_sc_magnet	3.0E+07	1	3.0E+07	3.0E+07		1.6E+06			
mttf_small_magnet	1.0E+07	1	1.0E+07	**3.4E+07**	3.4E+07				
mttf_PM_magnet	1.0E+07	1	1.0E+07	1.0E+07					
mttf_PS_corrector	4.0E+05	1	4.0E+05	**1.1E+06**	4.3E+05	1.8E+05	1.1E+05	1.1E+06	
mttf_PS	2.0E+05	5	1.0E+06	**1.1E+06**	4.3E+05	1.8E+05	1.1E+05	1.1E+06	4.0E+04
mttf_kicker	1.0E+05	1	1.0E+05	1.0E+05	1.0E+05				
mttf_kickpulser	7.0E+03	5	3.5E+04	7.0E+03	6.6E+03				
mttf_modulator	5.0E+04	1	5.0E+04	5.0E+04	6.4E+04				
mttf_dr_klystron	3.0E+04	1	3.0E+04	3.0E+04					
mttf_mb_klystron	4.0E+04	1	4.0E+04	4.0E+04	5.0E+04				
mttf_DRFS_klystron	1.2E+05	1	1.2E+05	1.2E+05					1.7E+05
mttf_X_klystron	2.5E+04	1	2.5E+04	2.5E+04					
mttf_cavity	1.0E+08	1	1.0E+08	1.0E+08					
mttf_coupler_intlk	1.0E+06	5	5.0E+06	1.0E+06	9.6E+04				
mttf_coupler_intlk_electronics	1.0E+06	1	1.0E+06	1.0E+06	9.6E+04				
mttf_mover	5.0E+05	1	5.0E+05	5.0E+05	5.1E+05				
mttf_VacP	1.0E+07	1	1.0E+07	1.0E+07	3.8E+06				
mttf_VacP_power_supply	1.0E+05	1	1.0E+05	1.0E+05					
mttf_valve	1.0E+06	1	1.0E+06	1.0E+06	1.0E+06				
mttf_vac_valve_controller	1.9E+05	1	1.9E+05	1.9E+05	1.9E+05				
mttf_fs	2.5E+05	10	2.5E+06	2.5E+05	2.2E+05				
mttf_pulsed_cable	2.0E+05	1	2.0E+05	2.0E+05					
mttf_xfrmr	2.0E+05	1	2.0E+05	2.0E+05					
mttf_waterpump	1.2E+05	1	1.2E+05	1.2E+05	1.2E+05	1.3E+05			
mttf_water_instr	3.0E+04	10	3.0E+05	**1.3E+05**	3.0E+04	1.3E+05			
mttf_elec_small	3.6E+05	1	3.6E+05	**1.6E+06**	3.6E+05				1.6E+06
mttf_elec_big	3.6E+05	1	3.6E+05	**1.6E+06**	3.6E+05			6.7E+05	1.6E+06
mttf_vac_mech_device	1.0E+05	5	5.0E+05	1.0E+05					
mttf_laser_wire	2.0E+04	1	2.0E+04	2.0E+04					
mttf_wire_scanner	1.0E+05	1	1.0E+05	1.0E+05					
mttf_klys_preamp	1.0E+05	1	1.0E+05	1.0E+05					
mttf_vacG_controller	1.0E+05	1	1.0E+05	**4.7E+05**	4.7E+05				
mttf_cavity_tuner	1.0E+06	1	1.0E+06	1.0E+06	5.1E+05				
mttf_cavity_piezo_tuner	5.0E+05	1	5.0E+05	5.0E+05					
mttf_power_coupler	1.0E+07	1	1.0E+07	1.0E+07					
mttf_SLED	1.0E+05	1	1.0E+05	1.0E+05					
mttf_cryo_leak	1.0E+05	1	1.0E+05	1.0E+05					
mttf_JT_valve	3.0E+05	1	3.0E+05	3.0E+05					
mttf_cryo_big_prob	1.0E+07	1	1.0E+07	1.0E+07					
mttf_target	4.4E+04	1	4.4E+04	4.4E+04					
simulations mttf_MPS_region	5.0E+03	1	5.0E+03	**3.0E+04**	5.0E+03				3.0E+04

For each item, the table gives the MTBFs used in the SB2009 simulations along with MTBFs from operating experience on four accelerators. The color codes indicate the extent of any gap between operating experience and the MTBFs used in the simulations. The color key is:

Green: MTBF has already been achieved
Yellow: MTBF is up to a factor 3 lower than used in the simulations
Gold: MTBF is up to a factor 10 lower than used in the simulations
Red: MTBF is more than a factor 10 lower than used in the simulation
White: no comparative data

Changes made to the Starting MTBFs from RDR to SB2009 are as follows:

Magnet Power Supplies
Lines 4, 9, and 10 were revised based on operations availability data from the Advanced Photon Source at Argonne, where an MTBF of 500,000hrs has been achieved over 30,000 operating hours of the ~1600 power converters in the APS storage ring. As an aside, it should be noted that there is no built-in redundancy in the APS power converters.

Superconducting Magnets
Line 6 was revised based on updated operations availability data from Fermilab.

Water Instrumentation
Line 31 was revised based on updated operations availability data from the Fermilab.

Electrical Utilities
Lines 32 and 33 were revised using data in the IEEE Gold Book, that includes surveys in-the-field failure rates for electrical utility equipment.

Machine Protection System
Line 47 was revised based on operations availability data from the Advanced Photon Source at Argonne.

Components for the Klystron Cluster and Distributed RF Systems
MTBF numbers for the new components in the two HLRF schemes are described earlier in this document.

4.4. Ingredients for meeting the required MTBFs
Irrespective of the exact optimization of Final MTBFs or the exact details of the model, it is clear there will be many challenges to building the ILC accelerator that can operate with the desired availability. There are many ingredients that must be considered that go beyond availability considerations in the design of the overall machine and of individual components. Two such ingredients are the role of Quality Assurance and Preventative Maintenance.

4.4.1. Quality Assurance / Quality Control
There are numerous examples where good reliable designs have been compromised in application by poor attention to detail in the implementation. Similarly, there are many examples of the benefits of effective quality control. Since almost every component in the ILC accelerator is duplicated many times, consistent and well controlled processes will be essential.

4.4.2. Preventative Maintenance
The fundamental principle of Preventative Maintenance is to take advantage of the scheduled down times in order to increase the availability during the scheduled operating periods by performing work that preemptively reduces the likelihood of equipment failures. Preventative Maintenance activities are planned in advance and carefully controlled in their implementation. The types of tasks considered

Preventative Maintenance take a wide range, including servicing of electrical switchgear and mechanical pumps, to replacement of water hoses when they reach 80% of their nominal useful life, to preemptive 'stress testing' of equipment during a shutdown in order to induce a failure in the weakest units so they can be removed from service.

5. Simulations

The simulation results described above were done with the program AVAILSIM developed by Tom Himel over the past several years. It is a Monte Carlo of the ILC that randomly breaks components, checks whether the set of presently broken components prevent the ILC from generating luminosity and if so, schedules repairs. A description of this program was given at the 2007 particle accelerator conference[i] and some of its features are described below.

Each component fails at a random time with an exponential distribution determined by its MTBF. When a component fails, the accelerator is degraded in some fashion. A klystron failure in the main linac simply reduces the energy overhead. The accelerator keeps running until this overhead is reduced to zero. Similarly there are 21 DR kickers where only 20 are needed so only the second failure causes downtime. Some components such as most magnet power supplies cause an immediate downtime for their repair.

Each component can be specified as hot swappable (meaning it can be replaced without further degrading the accelerator); repairable without accessing the accelerator tunnel, or repairable with an access to the accelerator tunnel. A klystron that is not in the accelerator tunnel is an example of a hot swappable device.

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, the simulation determines which parameter (e.g. e- linac energy overhead or e+ DR extraction kicker strength or luminosity) was degraded too much, and plans to fix things that degrade that parameter. Based on the required repairs, it calculates how long the downtime must be to repair the necessary items. It then schedules other items for repair, allowing the downtime to be extended by as much as 50 to 100%. Some other issues must also be taken into account:

• If an access to the accelerator tunnel is required, one hour is allowed for prompt radiation to decay before entry. One hour is also allowed for locking up, turning on and standardizing power supplies.

• The number of people in the accelerator tunnel can be limited to minimize the chaos of tracking who is in the tunnel. For the present work, we have not limited the number of people.

We have estimated that number and it is not unreasonably large, but we still need to use the simulation to limit that number and see how it affects the results.

• The accelerator runs a total of 9 months a year. There are two one month shutdowns for major maintenance and upgrades each year. In addition, every two weeks there is a 24 hour scheduled downtime to perform preventive maintenance. This includes work which is explicitly simulated like replacing DRFS klystrons which have died in the previous two weeks and work which is not explicitly simulated but which may be need to attain the long MTBFs like replacing pumps which are vibrating because their bearings are going bad.

• Recovery from a downtime is not instantaneous. Things break when you turn them off for the downtime and the problems are only discovered when they are turned back on. People make mistakes in hardware and software improvements and these must be found and corrected. Temperatures change and the ground moves so the beam must be re-tuned. Rather than trying to model downtime recovery procedures in detail, AVAILSIM simply assumes that the time it takes to get good beam out of a region of the accelerator is on average proportional to the time that region was without beam. The constants of proportionality used for each region typically were 10%, except for the DRs and interaction region, for which 20% was used and simple transport lines for which 5% was used. Altogether, this recovery time on average slightly more than doubles the downtime due to a repair.

The accelerator we simulated is not exactly the SB2009 design but we expect it is close enough that the simulation results are meaningful. We will continue our studies and add the remaining SB2009 changes in the near future. The simulated ILC had the following features:

- The linac could be in 1 or 2 tunnels
- Low power (half number of RDR bunches and RF power)
- Three RF systems were simulated for the main linac: RDR, KlyClus, and DRFS. Short linacs (injectors etc.) were always simulated with the RDR RF scheme.
- Two 6 km DRs in same tunnel near the IR. (Going to 3 km would decrease the component count and slightly improve the availability for all simulation runs.)
- RTML transport lines in linac tunnels
- Injectors in their own separate tunnels. (Putting them in the BDS will probably slightly decrease the availability for all simulation runs.)
- The e+ source uses an undulator at end of linac
- There is an e+ Keep Alive Source of adequate intensity to do machine development and to tune up during a downtime recovery.

- Injectors, RTML turn-around, DRs and BDS have all power supplies and controls accessible with beam on. (Pre-RDR 1 vs. 2 tunnel studies had these inaccessible for 1 tunnel.)

6. Outstanding issues and concerns

Availability studies, generally, require input on technology, system layout and operations and maintenance models. It is critically important, therefore, to make sure that the results of these studies are well documented and accessible to provide guidance for ongoing and future R & D activities.

The following outstanding availability issues and concerns have been identified. These are important for both the Reference Design and the proposed baseline.

- Integration with sub-system designers

  Component MTBF and MTTR performance requirements have been assigned and tabulated as a by-product of the availability analysis. In some cases, the performance criteria needed represent an advance beyond that routinely achieved in large colliders. In this case, industrial best-practice guidelines or light-source performance has been assumed. For both of these, added design effort, compared to large colliders, has to be applied.

- Estimating impact on cost

  Improved MTBF performance has a cost increment. In simple terms, higher quality off-the-shelf components, with adequate reliability performance, will cost more.

- Improving recovery times

  The simulation includes a 'luminosity recovery process' model, whereby the time to recover full-spec performance in each subsystem is proportional to the time without beam, (due, for example, to an upstream component failure). Completed simulation results indicated this to have a substantial impact, amounting to about half of the total simulated downtime. The proportionality constant used is based on observations done at several facilities, and some degree of variability is seen. A consistent approach to improving system availability includes effort on both component MTBF and the associated recovery process.

- Safety – impact on Operations and Maintenance model

  Each intervention for maintenance, including preventive activities and incident response, must be subject to a fully integrated safety analysis. The simulation includes several aspects, such as radiation cool-down wait-times and an allowance for proper preparation and entry procedures. It is reasonable to expect additional constraints.

- Accessibility – equipment transport

  For repairs within the beamline enclosure, the availability model includes an estimate of the time to move equipment from a storage area to the access-way. With more detailed underground equipment layout, the model will be updated to provide a more realistic estimate of transit time.

- Radiation effects

  Radiation effects may reduce electronic component MTBF substantially. This is important because the new baseline layout does not provide nearby support equipment enclosure space for the main linac and RTML. Several accelerator labs have experience with electronics located near the beamline, notably the PEPII B-Factory and (soon) the LHC. The availability simulation assumes that electronics is protected from radiation and does not take into account possible reduced performance.

- Initial commissioning

  None of the studies include an estimate of initial commissioning effects.

- Staffing levels

  The model shows, correctly, the balance between component failure rate and total time to replace or repair. An important ingredient of that balance is the estimate of the staffing level so that maintenance activities can proceed in parallel. Due to the statistical nature of component failure, the simulation can provide a good estimate of the required staffing levels.

- MTBF data

  The basis for the estimate is the reliability data collected from labs and industrial literature. In many cases, representative estimates are difficult to obtain, so the analysis has been based on data from only one institute, or is simply an estimate. There are several reasons why valid MTBF estimates are difficult to obtain, notably the tendency for organizations to group failures according to who in their organization is responsible for repairs associated with a given subsystem.

---

[i] Proceedings of PAC 2009, Albuquerque, New Mexico, p. 1966-9