LC availability Simulation done for the LC comparison task force

Tom Himel
SLAC

Committee Members

- Tom Himel (chair)
- Paul Czarapata
- Helen Edwards
- Markus Huening (added since last meeting)
- Nan Phinney
- Marc Ross

Contents

- Our approach to estimating an accelerator's availability
- The simulation
- The state of the input data
- Very preliminary results
- Plans

Overall plan of Attack

- Write a simulation that given the MTBFs, MTTRs, numbers and redundancies of components, and access requirements for repair can calculate average availability and the integrated luminosity per year. Luminosity will mostly be either design or zero in this simulation.
- Collect data on MTBFs and MTBRs of components in existing machines to guide our budgeting process
- Make up a reasonable set of MTBFs that give a reasonable overall availability.
- Iterate as many times as we have time for (probably once during this task force) to minimize the overall cost of the LC while maintaining the goal availability

Why a simulation?

- We chose to go with a simulation instead of a spreadsheet calculation for the following reasons:
  - Including tuning and recovery times in a spreadsheet calculation is difficult.
  - Fixing many things at once (during an access) is also difficult to put in a simple spreadsheet formula.
  - If later, one wants to more carefully model luminosity degradation on recovery from downtimes a simulation is simpler
  - A disadvantage of a simulation is its use of random numbers so one needs high enough statistics go get a meaningful answer. This is particularly a concern if one wants to compare two slightly different cases. Random number seeds will be handled in a way to allow meaningful comparisons of similar cases.

Simulation overview

- About 1100 lines of Matlab code
- Reads an Excel spreadsheet with data that describes the accelerator
- Lets things break according to their MTBFs
- Evaluates how the accelerator performance degrades
- Schedules repairs when performance is too poor
- Does crude accounting of retuning time after a repair. Accrues opportunistic machine development time.
- Outputs a comma delimited file for import to Excel that has the downtime caused by each device.

Accelerator parameters

- The simulation keeps track of various accelerator parameters which degrade as components break. Examples:
  - e- linac energy overhead
  - e+ damping ring extraction kicker strength
  - e- damping ring RF voltage
  - Luminosity
- If a parameter gets below its minimum allowed value, the LC is declared broken and repairs are scheduled.

Actual parameter input sheet

name	design	minimum
luminosity	1.00E+34	5.00E+33
e- energy overhead	5000	0
e+ energy overhead	5000	0
e- DR RF HV	54	49.49955
e+ DR RF HV	54	49.49955
e- DR inj kick	0.63003	0.6
e- DR ext kick	0.63003	0.6
e+ DR inj kick	0.63003	0.6
e+ DR ext kick	0.63003	0.6

2% energy overhead

11 of 12 RF cavities

20 of 21 working inj and extraction kickers

Few lines of Actual Component Input Sheet

rank in subsys	component name	subsys/segment	region	problem name	quantity	parameter effected	add/ mult	degradation	MTBF	MTTR	Still broken after repair	access needed ?	n repair people	rand seed
1	beamline component	beamline	e+ comp	broken	1	luminosity	mult	0	3.00E+04	8		1	2	9
1	off beamline component	beamline	e+ comp	broken	1	luminosity	mult	0	1.50E+03	1		0	2	10
2	Quads	beamline	e- linac	broken	253	luminosity	mult	0.00	1.00E+08	2	quad or	0	1	22
3	Corrs	beamline	e- linac	broken	379	luminosity	mult	0.00	1.00E+08	0.5	quad or	0	1	23
3	quad or corr	beamline	e- linac	retuned	632	luminosity	mult	0.99	1.00E+50	2		0	2	24
4	Power supplies - bend	beamline	e- linac	broken	3	luminosity	mult	0.00	2.00E+05	2		0	2	25
4	Power supplies - quad	beamline	e- linac	broken	253	luminosity	mult	0.00	2.00E+05	2	quad or	0	2	219
4	Power supplies - corr	beamline	e- linac	broken	379	luminosity	mult	0.00	2.00E+05	0.5	quad or	0	1	26
5	PS controller - bend	beamline	e- linac	broken	3	luminosity	mult	0.00	5.00E+05	1		0	1	27
9	Water pumps	beamline	e- linac	broken	3	luminosity	mult	0.00	1.20E+05	4		1	2	32
10	Water	beamline	e- linac	broken	3	luminosity	mult	0.00	3.00E+04	2		1	2	33
11	Flow Switch	beamline	e- linac	broken	6	luminosity	mult	0.00	2.30E+05	1		1	2	34
1	Cavities	cavity	e- linac	degrade	7152	e- energy ov	add	-12.00	1.00E+08	336		1	2	35
1.1	Cavities	cavity	e- linac	broken	7152	e- energy ov	add	-36.33	1.00E+08	336		1	4	36
2	Cavity tuner	cavity	e- linac	broken	7152	e- energy ov	add	-36.33	5.00E+05	336		1	4	37
3	Cavity piezo tuner	cavity	e- linac	broken	7152	e- energy ov	add	-12.00	5.00E+05	2	power c	1	4	38
4	LLRF	cavity	e- linac	broken	7152	e- energy ov	add	-36.33	1.00E+05	1		0	1	39
1	power coupler	coupler	e- linac	degrade	7152	e- energy ov	add	-288.00	1.00E+07	2	power c	1	2	40
1.1	power coupler	coupler	e- linac	broken	7152	e- energy ov	add	-871.92	1.00E+07	2	power c	1	2	41
1.1	power coupler disc	coupler	e- linac	disc	7152	e- energy ov	add	-36.33	1.00E+50	336		1	4	42
1	Klystrons	klystron	e- linac	broken	293	e- energy ov	add	-871.92	4.00E+04	8		-1	2	63
1.5	pulse transformers	klystron	e- linac	broken	293	e- energy ov	add	-871.92	1.00E+05	4		-1	2	222
2	Modulators	klystron	e- linac	broken	293	e- energy ov	add	-871.92	2.00E+04	2		0	2	64
2.5	pulse cables	klystron	e- linac	broken	293	e- energy ov	add	-871.92	1.00E+06	4		0	2	65
2	Quads	beamline	e- DR	broken	849	luminosity	mult	0.00	4.90E+06	8		1	2	77
3	Sextupoles	beamline	e- DR	broken	312	luminosity	mult	0.00	4.90E+06	8		1	2	78
4	Corrs	beamline	e- DR	broken	629	luminosity	mult	0.00	1.00E+08	0.5	quad or	0	2	79
5	Wigglers	beamline	e- DR	broken	90	luminosity	mult	0.00	4.90E+05	8		1	2	80

Component Sheet Columns Defined

- Rank in subsys: Not used by simulation. Helps sheet organization
- Component name e.g. klystron or modulator or bend magnet
- Subsys/segment: Not used by simulation. Helps sheet organization
- Region: Part of the accelerator, e.g. e- linac or e+ source
- Problem name: Allows for problems with different consequences
- Quantity: Total number of this device in this region
- Parameter effected: e.g. e- energy overhead or luminosity
- Add/mult: Affect on the parameter is additive or multiplicative
- Degradation: Amount the parameter is degraded for each broken device
- MTBF: Mean Time Between Failures (hours)
- MTTR: Mean Time To Repair (hours) not including time to access & recover
- Still broken after repair: Component name of what is still broken after a kludged repair is complete. E.g. corrector supply still needs to be fixed after temporary fix of steering around it. (Not implemented yet)
- Access needed: =1 means access to the accelerator tunnel is needed for a repair
  =0 means no access is needed, but accelerator is down during repair
  =-1 means component is hot swappable e.g. a klystron in 2 tunnel case.
- n repair people: number of people needed to repair a component (specialties are ignored)
- Rand seed: Allows use of same random numbers even when a new component line is added

The Ful Components Sheet

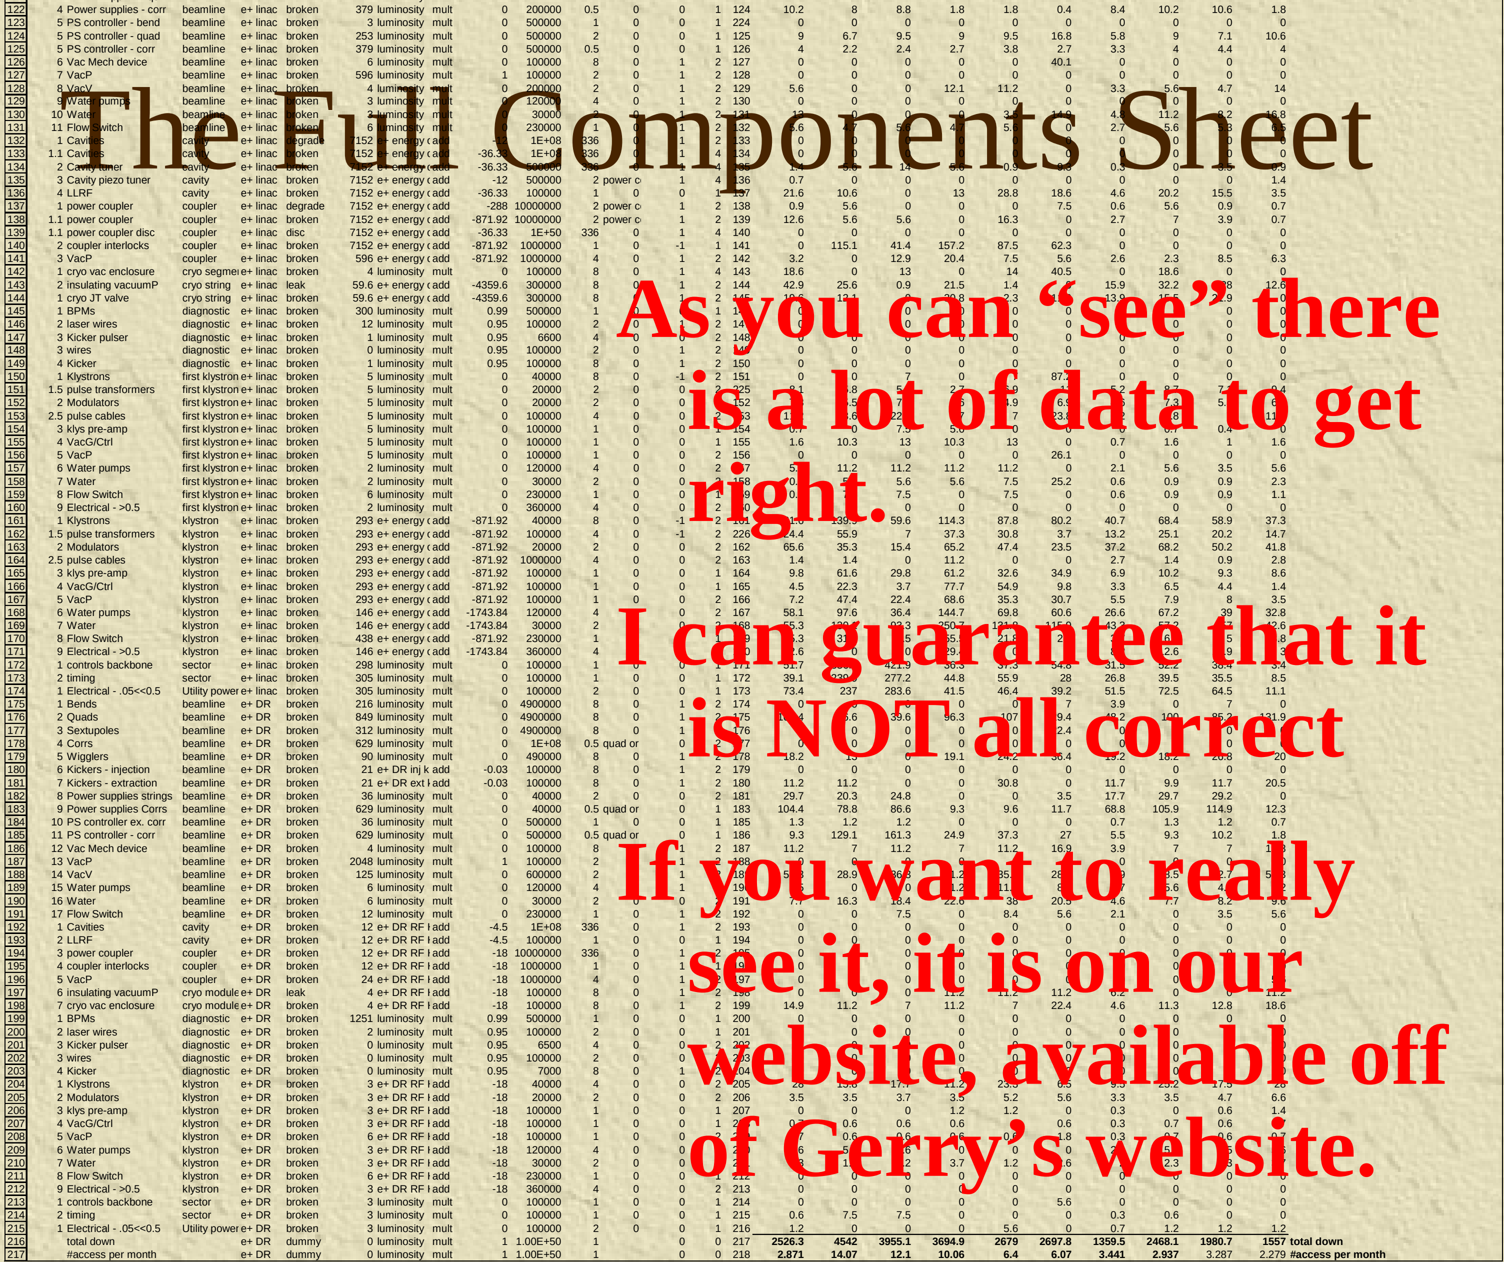

As you can "see" there is a lot of data to get right.

I can guarantee that it is NOT all correct

If you want to really see it, it is on our website, available off of Gerry's website.

Tom Himel

Lumped Availability Budgets

- Due to time constraints, only linacs and DRs were modeled in detail.
- Here are downtime budgets we assigned to other systems.
- Actually unavailabilities will be roughly double these numbers due to recovery times.
- Total unavail goal is 25%. 10% contingency is unbudgeted, leaving 15% to budget.

region	**warm nominal % downtime**	**cold nominal % downtime**
e- injector	0.48	0.36
e- compressor	0.48	0.36
e- BDS	0.48	0.36
e+ source	0.48	0.36
e+ PDR	0.48	1.00E-40
e+ compressor	0.48	0.36
e+ BDS	0.48	0.36
IP	0.48	0.36
cryo plant	1.00E-40	1
site power	0.5	0.5
Sum	**4.34**	**4.02**
2*Sum	**8.68**	**8.04**

Modeling breakdowns

- At initialization calculate the next time each component will break.
  - Average time is MTBF/(number of the component)
  - Throw a random number with an exponential distribution with the above average.
- When that time comes, degrade the corresponding parameter and calculate the next time one will break.
- If that parameter is too far degraded, then immediately plan and start repairs.
- All the downtime and recovery time from the repairs is accounted to the component which put us over the edge even though other components will also be repaired.
- To keep things simple, if something breaks when we are in the middle of repairs, it is just ignored. In real life, sometimes it would be noticed and fixed, and other times it would contribute to the recovery time.

Downtime planning

- This is without doubt the most complicated part of the simulation.
- Anyone who has participated in a downtime scheduling meeting will understand why.
- And computers don’t get a gestalt of the situation like humans can.
- It’s hopefully good enough

Downtime planning

1. Check each component to see if by itself it “breaks” the accelerator. If there is then fix all of those components. For example if 2 DR kickers are broken then they would both be scheduled for repair.
2. Next we handle the case where multiple components together break the accelerator. For example, broken klystrons and cavity tuners both reduce the energy, but only together do they reduce it enough to run out of energy headroom.
   - Loop through the components in order of increasing “bang for the buck” (improvement per repair hour)
   - Accumulate parameter degradation (e.g. linac energy overhead)
   - Repair all components after the degradation gets too big
   - This gets the accelerator fixed in the minimum possible time.

Downtime planning

3. The above repairs are scheduled taking into account the available number of repair people (set to 25 in the accelerator tunnel and 100 outside it) and the access times for the regions. These repairs are enough to make the accelerator work again. Note that the fact that repair people have specialties has been ignored to keep things simple.
4. Now schedule extra repairs .
   - In this case, we plan repairs which give the most “bang for the buck” first.
   - Keep scheduling things until it would extend the downtime by too much (set to 50% if no access is being done and a factor of two if there is an access)
   - Note there is no logic to decide which degraded parameter should be addressed first. Hence the simulation could choose to repair e- linac klystrons before e+ linac klystrons even though there happens to be less headroom in the e+ linac.

Regions and PPS zones

name	**ppszone**	**upstream**	**accesshours**	**recoverhours**	**tunetimefraction**	**start of beamline**
none	none	none	0	0	0	0
e- injector	e- injector	none	2	1	0.1	1
e- DR	e- linac	e- injector	3	1	0.2	0
e- compressor	e- linac	e- DR	3	1	0.1	0
e- linac	e- linac	e- compressor	3	1	0.1	0
e- BDS	e- BDS	e- linac	3	1	0.1	0
e+ source	e+ source	e- linac	2	1	0.1	1
e+ DR	e+ linac	e+ source	3	1	0.2	0
e+ compressor	e+ linac	e+ DR	3	1	0.1	0
e+ linac	e+ linac	e+ compressor	3	1	0.1	0
e+ BDS	e+ BDS	e+ linac	3	1	0.1	0
IP	IP	e+ BDS	1	1	0.2	0

- Two regions are in the same PPS zone if access to one of them means there can be no beam in the other. E.g. with the DR and linac in the same tunnel, they are in the same PPS zone.
- “Upstream” indicates the order the beam goes through the regions.
- “tunetimefraction” is used for the crude way we simulate recovery from a downtime. The time with no beam in a region times the tunetimefraction gives the time it takes to get good beam to the end of that region. At that time tuning can begin in the next downstream region.

Machine Development time Budget

- We budget % of time needed for MD in each region.
- Note conventional e+ source allows simultaneous MD in e.g. both DRs.
- Some MD will get done opportunistically while other parts of accel are down.
- Rest must be scheduled.

name	**%MD time warm**	**%MD time cold**
e- injector	1	1
e- DR	2	2
e- compressor	1	1
e- linac	1	1
e- BDS	1	1
e+ source	1	1
e+ PDR	1	0
e+ DR	2	2
e+ compressor	1	1
e+ linac	1	1
e+ BDS	1	1
IP	1	1
total conv e+ source	**8**	**7**
total wiggler e+ source	**13**	**12**

Modeling recovery

- It is assumed that everything goes exactly according to schedule and the exact MTTRs are used as repair times (no random numbers thrown for this).
- Hence the time work will be done in each region and PPS zone is easily determined.
- Working our way downstream from each injector, we see how long the region has been without beam, use tunetimefraction to calculate when beam will be at the end of the region and then go on to the next region.
- If the next region is still undergoing repairs that will continue for more than 2 hours, we assume opportunistic machine development will be done in the part of the accelerator with beam. This reduces the amount of scheduled MD time which must be done.

Modeling Long shutdowns

- This is very simple.
- Every 9 months (an input parameter) all devices which take more than 300 hours to repair are instantaneously repaired.
- That's it.
- Other devices are not repaired as some probably break during the long shutdown anyway.
- Recovery from a long shutdown is not modeled. This should be done by hand when estimating annual integrated luminosity. For a 3 month shutdown, assuming 1.5 months with no luminosity would do a good job of mimicking the probably 3 month slow recovery of luminosity.

Misc Parameter sheet

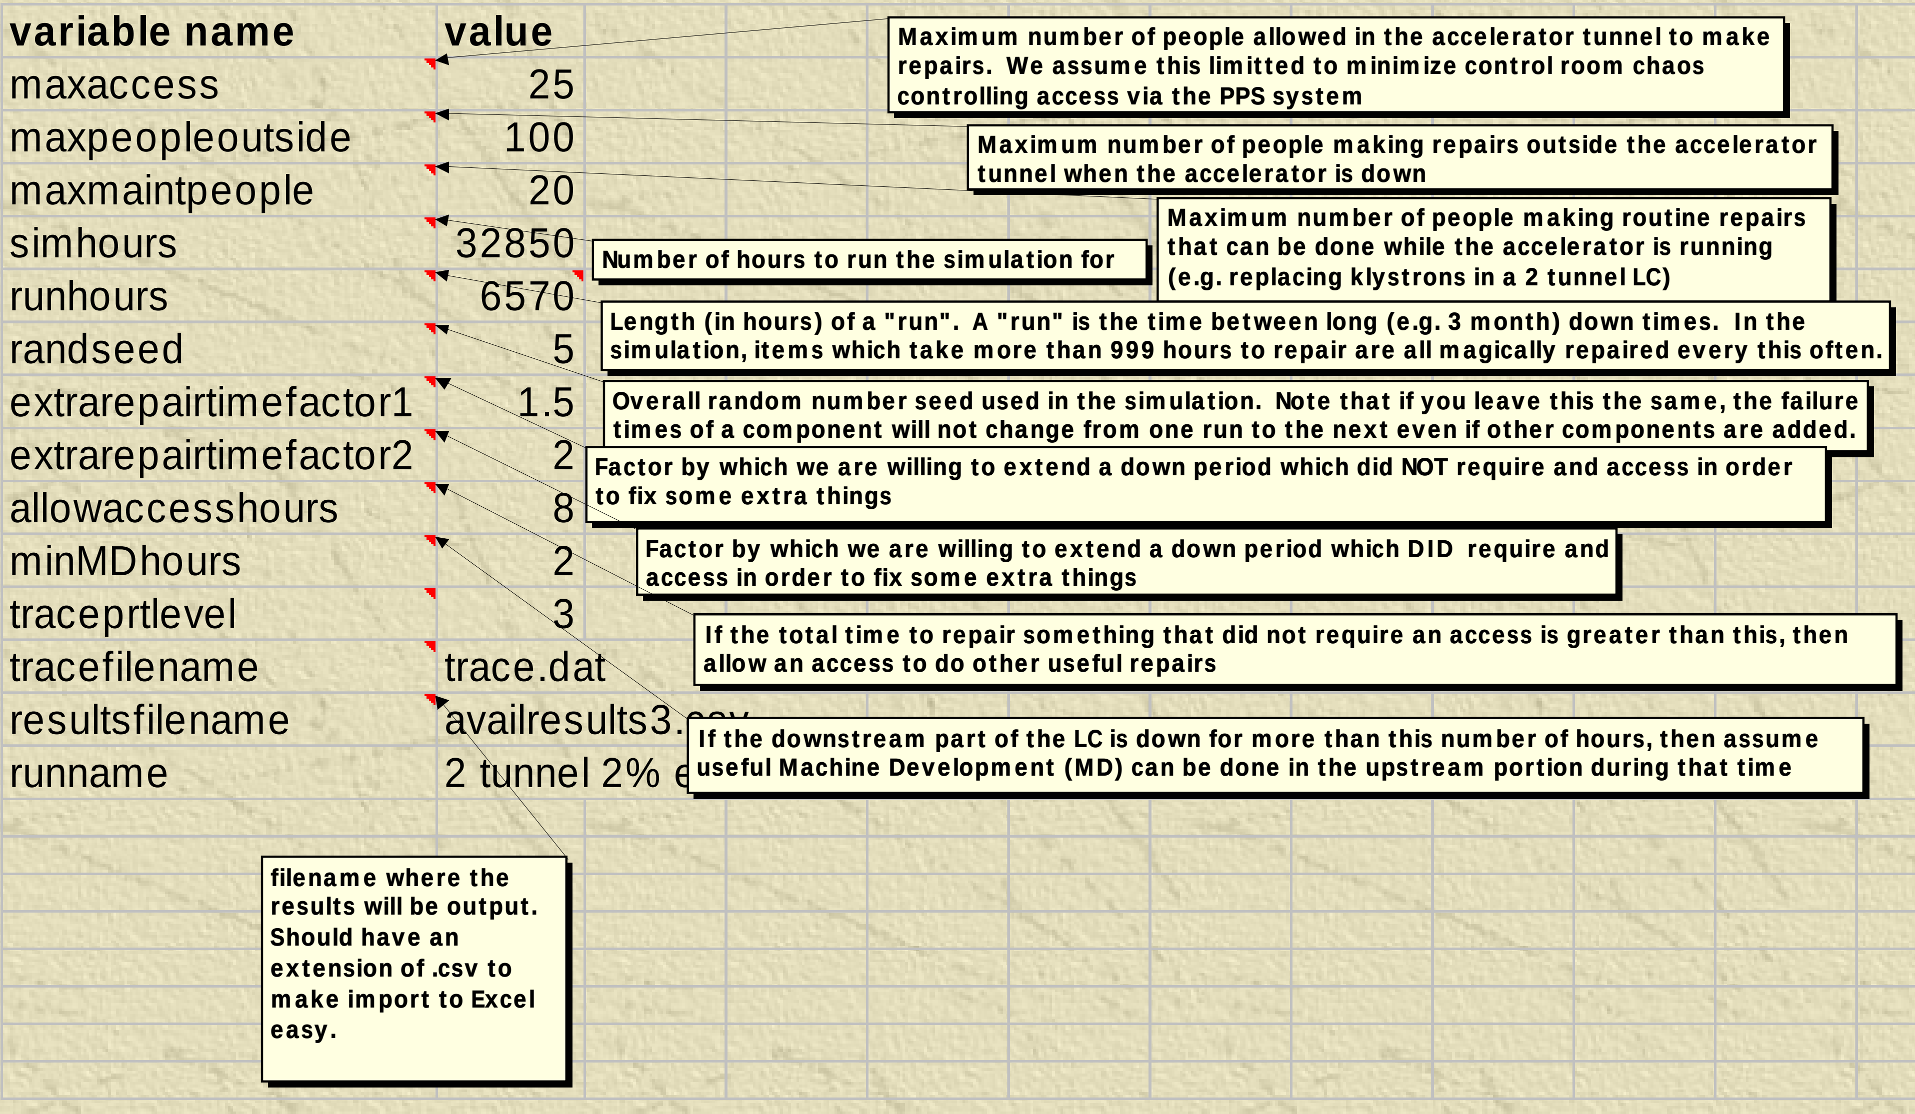

variable name	value
maxaccess	25
maxpeopleoutside	100
maxmaintpeople	20
simhours	32850
runhours	6570
randseed	5
extrarepairtimefactor1	1.5
extrarepairtimefactor2	2
allowaccesshours	8
minMDhours	2
traceprtlevel	3
tracefilename	trace.dat
resultsfilename	availresults3.
runname	2 tunnel 2% e

Maximum number of people allowed in the accelerator tunnel to make repairs. We assume this limitted to minimize control room chaos controlling access via the PPS system

Maximum number of people making repairs outside the accelerator tunnel when the accelerator is down

Maximum number of people making routine repairs that can be done while the accelerator is running (e.g. replacing klystrons in a 2 tunnel LC)

Number of hours to run the simulation for

Length (in hours) of a "run". A "run" is the time between long (e.g. 3 month) down times. In the simulation, items which take more than 999 hours to repair are all magically repaired every this often.

Overall random number seed used in the simulation. Note that if you leave this the same, the failure times of a component will not change from one run to the next even if other components are added.

Factor by which we are willing to extend a down period which did NOT require and access in order to fix some extra things

Factor by which we are willing to extend a down period which DID require and access in order to fix some extra things

If the total time to repair something that did not require an access is greater than this, then allow an access to do other useful repairs

If the downstream part of the LC is down for more than this number of hours, then assume useful Machine Development (MD) can be done in the upstream portion during that time

filename where the results will be output. Should have an extension of .csv to make import to Excel easy.

State of input data

- Only DRs and linacs are modeled in detail
- Have reasonably good parts counts
    - Some components could still be missing (definitely pulsed cables and 3 stub tuners)
    - Need to make clearer definitions of each component: e.g. are transducers considered part of the power supply, the controller, or as a separate part.
- MTTRs may be good to a factor of 2.
- Just learning the consequences when a part breaks was far more difficult then expected and still needs more work.
- Tuning model assumes there is a tune-up dump at the end of each region.

State of input data

- MTBFs are the big problem
  - In many cases, don’t have data from present accelerators to get a good starting value
  - Often just used 100,000 hours for no good reason.
  - Ideally we would know MTBF vs. cost for each component and minimize total cost subject to the constraint on overall availability – We knew upfront this wasn’t possible.
  - For a few which are real availability drivers we will look in detail. Others will just be guesses.
  - Even calculating the MTBF of a power supply with redundant regulators is difficult as common mode failures (e.g. water leak) must be considered.
- Remember: GIGO

All Results are Preliminary

- Many MTBFs are just defaulted to 100k hours
- Haven't budgeted enough of the unavailability to the non-DR, linac regions
- Almost NO optimization of the MTBFs has been done
  - Just took initial set of MTBFs, used the simulation to see which components made major contributions to the downtime and then increased their MTBFs.
- Haven't checked many sensitivities e.g. to the number of repair people.

Results can be useful anyway

- They highlight the components whose MTBFs must be increased.
- The general magnitude (and sign) of availability changes with major design changes should be reasonable.
  - At present have done too few runs to judge how sensitive these changes are too other inputs.
  - There are so many inputs, hopefully their errors average out.

Tom Himel

Very Preliminary Results

% time down	# access per month	# tunnels	energy overhead	MTBF fudge	special conditions
25.3	2.9	2	2%	1	DR magnet MTBF = 10*SLACs = 4.9e6 hours DR vacuum valve MTBF = 6e5 instead of nominal 1e5
45.4	14.1	1	2%	1	
39.6	12.1	1	4%	1	
36.9	10.1	1	2%	10	MTBFs of magnet power supplies, timing, controls backbone, and electrical distribution increased by a factor of 10.
26.8	6.4	1	4%	10	
27.0	6.1	1	4%	10	A different random number seed to see how precise the results are. The total downtime changes very little. The components vary by around 0.25%.
13.6	3.4	2	2%	1	reduces the tuning time by a factor of 10 to give and idea of how important that tuning time is.
24.7	2.9	2	2%	1	DR is separate tunnel. A surprisingly small effect
19.8	3.3	2	2%	1	Conventional e+ source: Just assumed e+ could be produced with no beam in e- linac. No changes were made to the reliability of parts in the e+ source itself. This illustrates the penalty of not being able to tune up (or keep tuned) the positrons when something in the electron system is down
15.6	2.3	2	4%	10	

Warm MTBF budgets

run number	% time down incl forced MD	% time fully up integrating lum or sched MD	% time integrating lum	% time scheduled MD	% time actual opportunistic MD	% time useless down	MTBF improvement factor	MTBF before	MTBF now	devices
warm1	31.8	68.2	59.5	8.7	4.3	27.5				initial run with nominal MTBFs
							10	1.00E+06	1.00E+07	all water cooled magnets
							5	5.00E+04	2.50E+05	DR Q movers
							10	2.00E+05	2.00E+06	large Power Supplies
							10	1.00E+05	1.00E+06	electronics modules esp PS controllers
							3	2.50E+05	7.50E+05	flowswitches
							50	2.50E+04	1.25E+06	first klystrons and related hardware (should be done with redundancy)
warm2	20.5	79.5	70.6	8.9	4.08	16.4	5	3.60E+05	1.80E+06	small electrical < 0.5 MW
							50	1.00E+06	5.00E+07	all water cooled magnets (total improvement)
							5	5.00E+04	2.50E+05	linac RF movers
							20	1.00E+05	2.00E+06	electronics modules esp PS controllers (total improvement)
warm3	16.6	83.4	73.5	9.9	3.1	13.5	10	5.00E+04	5.00E+05	linac vacuum mechanical devices
warm4	12	88	85.1	2.9	5.15	6.87				warm3 but with conventional e+ target
warm5	52.1	47.9	24.6	23	2.64	49.4				warm3 but 0.5x all MTBFs, 2x all tunetimefraction, 2x all MD times
warm6	30.8	69.2	66.8	2.4	13.6	17.2				warm 5 but with conventional e+ target

Cold MTBF budgets

run number	% time down incl forced MD	% time fully up integrating lum or sched MD	% time integrating lum	% time scheduled MD	% time actual opportunistic MD	% time useless down	MTBF improvement factor	MTBF before	MTBF now	devices
cold1	30.9	69.1	58.6	11	1.49	29.4				initial run with nominal MTBFs
							10	1.00E+06	1.00E+07	all water cooled magnets
							3	5.00E+04	1.50E+05	vacuum valve controllers
							5	2.00E+05	1.00E+06	large Power Supplies
							10	1.00E+05	1.00E+06	electronics modules esp PS controllers
							3	1.00E+05	3.00E+05	DR coupler interlocks
							5	1.00E+05	5.00E+05	DR VACP
							5	3.00E+05	1.50E+06	vacuum insulating vacuum leaks
										energy overhead increase from 2 to 3%
cold2	18	82	71.1	11	1.09	16.9	3	3.60E+05	1.08E+06	small electrical < 0.5 MW
							4	5.00E+04	2.00E+05	linac vacuum mechanical device
										energy overhead increased to 4% (total improvement)
							30	1.00E+06	3.00E+07	all water cooled magnets (total improvement)
							10	2.50E+04	2.50E+05	first klystrons and related hardware (should be done with redundancy)
cold3	16.6	83.4	72.5	11	1.11	15.5	3	3.00E+05	9.00E+05	cryo JT valve
cold4	35.2	64.8	54.2	11	1.41	33.8				cold 3 but with 1 tunnel
cold5										cold4 but with improved MTBFs. Got tired, so not done yet.

Comments

- Many parts means each must be quite reliable.
- Of initial 30% unavail, 10% is for lumped systems, 20% for linac+DR. Hence need factor of 4 overall improvement for them.
- Wiggler e+ source really hurts availability for 3 reasons
  - MD can't be done in parallel
  - Less MD can be done during recovery
  - Sequential recovery is slower
- Cold and warm are remarkably similar
- Some serious engineering will be needed to attain the required MTBFs.

Plans

- Cleanup input deck more
- Rerun simulations – very similar results expected
- Write it up
- Celebrate