PEP-II Hardware Reliability

C. W. Allen, W. Colocho, R. Erickson, M. Stanek

Stanford Linear Accelerator Center, Stanford, CA[1]

***Abstract*--Hardware reliability takes on special importance in large accelerator facilities intended to work as factories; i.e., when they are expected to deliver design performance for extended periods of time. The PEP-II "B-Factory" at SLAC is such a facility. In this paper, we summarize PEP-II reliability statistics from the first four years of production running. The four running periods extended from January 12 through October 31, 2000, from February 4, 2001 through June 30, 2002, from November 15, 2002 through June 30, 2003, and from September 9, 2003 through July 31, 2004. These four periods are designated Runs 1, 2, 3, and 4 in the discussion and tables presented in the paper. The first four runs encompassed 30,359 hours. During this time, PEP-II was delivering luminosity to the BaBar detector 57.9 percent of the time. In addition, 5.3 percent of the time was used for scheduled dedicated machine development work, and 4.5 percent was scheduled off for maintenance, installations, or safety checks. Injection and tuning accounted for 19.9 percent. The remaining 12.4 percent was lost due to malfunctions. During this time period, a total of 9701 malfunctions were reported, but most did not interrupt the running program. The unscheduled down time, a total of 3883 hours, was attributed to 1724 of these malfunctions. Mean Time to Fail (MTTF) and Mean Time to Repair (MTTR) are presented for each of the major subsystems, and long-term availability trends are discussed.**

I. Introduction

THE PEP-II B-FACTORY at SLAC, which consists of two intersecting storage rings filled with electrons and positrons from the SLAC two-mile linear accelerator, was designed to explore rare processes associated with the production and decay of b quarks produced in e^+e^- collisions. These studies require continuous operation for extended periods of time at very high luminosity in order to collect the requisite large data sets. Hence, hardware reliability takes on special importance. In this paper, we summarize the reliability statistics from the first four years of PEP-II production running. The four running periods extended from January 12 through October 31, 2000, from February 4, 2001 through June 30, 2002, from November 15, 2002 through June 30, 2003, and from September 9, 2003 through July 31, 2004. These four periods are designated Runs 1, 2, 3, and 4 in the discussion and tables that follow. Annual holiday shut-down periods are not counted in the statistical averages; nor are the long down time periods between the runs. The first four runs encompassed 30,359 hours. During this time, PEP-II was delivering luminosity to the BaBar experimental detector 57.9 percent of the time. In addition, 5.3 percent of the time was used for scheduled dedicated machine development work, and 4.5 percent was scheduled off for maintenance, installations, or safety checks. Injection and tuning accounted for 19.9 percent. The remaining 12.4 percent was lost due to malfunctions. During this time period, a total of 9701 malfunctions were reported, but most did not interrupt the running program. The unscheduled down time, a total of 3883 hours, was attributed to 1724 of these malfunctions.

II. Methodology

The data presented here are summarized from a database of all hardware malfunctions observed during the operation of the PEP-II facility. This database, known as ARTEMIS, is used to track problems, beginning with the initial symptoms identified by operators, and including information about repair tasks assigned, resolution of the initial problems, close-out issues as appropriate, and any adverse impact of each malfunction on the scheduled accelerator program. In general, invasive maintenance activities are not scheduled on a routine basis at the PEP-II facility. Instead, the ARTEMIS database, along with other organizational tools, is used to generate lists of pending maintenance tasks. When a malfunction interrupts the accelerator program, a maintenance crew is typically dispatched to make repairs, and if appropriate, other crews are dispatched to carry out tasks from the previously prepared lists. This "opportunistic" approach to maintenance has worked very well at SLAC [1].

In the discussion that follows, malfunctions that resulted in lost beam time are referred to as "events". Mean Time to Fail (MTTF), Mean Time to Repair (MTTR), and availability are defined as follows:

MTTF = Scheduled beam time / events.
MTTR = Unscheduled down time / events.
Availability = 1 – Unscheduled down time / Scheduled beam time.

[1] This work was supported by Department of Energy contract DE-AC02-76SF00515.

The recovery time following stored beam aborts is not counted as downtime unless the program is stopped to repair something and a database entry is generated. Otherwise, the recovery from an abort is counted as injection and tuning time.

III. Reliability Data

Table I summarizes the data set for the four combined runs, sorted by subsystem (power supplies, magnets, and so on), and again by area (injector, north damping ring, and so on).

TABLE I
HARDWARE RELIABILITY STATISTICS

	Scheduled Operating Hours (SOH)		30332.2
MTTF	17.59		
MTTR	2.25	Total Events	1724
Availability	87%	Total Hrs down	3883.0
Scheduled Operating Hours = 31,920 - Scheduled Outage			
Scheduled Outage =	1587.8		

Subsystem Data

System	Events	Hrs Down	% SOH	MTTF	MTTR
Power Supplies	565	1010.8	3.33%	53.7	1.8
Magnets	49	226.2	0.75%	619.0	4.6
RF	308	632.4	2.08%	98.5	2.1
Vacuum	76	407.3	1.34%	399.1	5.4
Utilities	198	617.3	2.04%	153.2	3.1
Cryogenics	16	81.2	0.27%	1895.8	5.1
Controls	375	577.7	1.90%	80.9	1.5
Safety	4	10.9	0.04%	7583.1	2.7
Other	133	319.2	1.05%	228.1	2.4
Totals	**1724**	**3883.0**		**17.59**	**2.25**

Area Data

Area	Events	Hrs Down	% SOH	MTTF	MTTR
CID	10	19.2	0.06%	3033.2	1.9
Sect. 0/1	84	190.7	0.63%	361.1	2.3
WTA	1	0.2	0.00%	30332.2	0.2
NDR	113	276.0	0.91%	268.4	2.4
SDR	111	278.7	0.92%	273.3	2.5
DRIP	5	29.4	0.10%	6066.4	5.9
Linac	159	221.3	0.73%	190.8	1.4
Positron	82	136.5	0.45%	369.9	1.7
PEPII-Inj	55	84.3	0.28%	551.5	1.5
BSY	18	24.6	0.08%	1685.1	1.4
HER	526	1045.2	3.45%	57.7	2.0
LER	306	837.2	2.76%	99.1	2.7
Near IR	11	27.8	0.09%	2757.5	2.5
MCC	37	53.4	0.18%	819.8	1.4
ESA	5	9.3	0.03%	6066.4	1.9
FFTB	4	2.1	0.01%	7583.1	0.5
Other	197	647.1	2.13%	154.0	3.3
	1724	**3883.0**		**17.59**	**2.25**

The data illustrate that magnets and vacuum systems have relatively long MTTF values, but also have relatively long MTTRs. These systems tend to be inherently robust, but when they fail, a tunnel entry is required to make repairs, and single events can sometimes require many hours to repair. In contrast, power supplies and controls, which typically have large numbers of active components, fail more frequently (shorter MTTF), but are also repaired more quickly (shorter MTTR). These systems are typically engineered with replaceable modular components, and are usually located outside the tunnel enclosures.

Weekly averages of MTTF, MTTR, and availability are plotted in Figure 1. In general, the availability improved slowly over the course of each run as the frequency of failures declined. In Run 2, the availability reached a plateau in the latter part of 2001. During the first part of the run, the peak luminosity improved significantly as efforts were made to maximize the stored currents, often by running near the technical limits of the hardware. In the latter half of Run 2, the machine was allowed to run steadily for long periods, and technically risky machine parameters were avoided. This conservative mode of operation yielded very good hardware availability, but little progress was made during this period in raising the peak luminosity.

For the first few weeks of Run 3, the MTTF and availability were low. This was due partly to frequent random malfunctions that typically appear during the first few days following a long down time. This time, however, a new abort kicker system featuring a significantly faster rise time was commissioned and required frequent interventions to debug. Then, in late December, 2002, a thunder storm disrupted the electric power, causing numerous problems from which the machine never fully recovered before the holiday shut down period.

Beginning in 2003, PEP-II ran well, with availability typically above 90 percent. As the run progressed, the stored currents were increased until the machine performance was limited by vacuum chamber failures due to beam-induced heating. The deep dips in the availability plot in 2003 were due to these vacuum chamber failures

This plot is somewhat misleading in accounting for the full impact of vacuum system failures. Sometimes vacuum failures first appear as very small leaks and can be patched in place without venting or disassembling the vacuum system. On a few occasions, small leaks appeared at night or on weekends and were patched quickly, and invasive or time-consuming repairs were deferred until the next regular work day. In these cases, only the initial repair time is counted as unscheduled down time; the major repair work that came later was counted as scheduled off time.

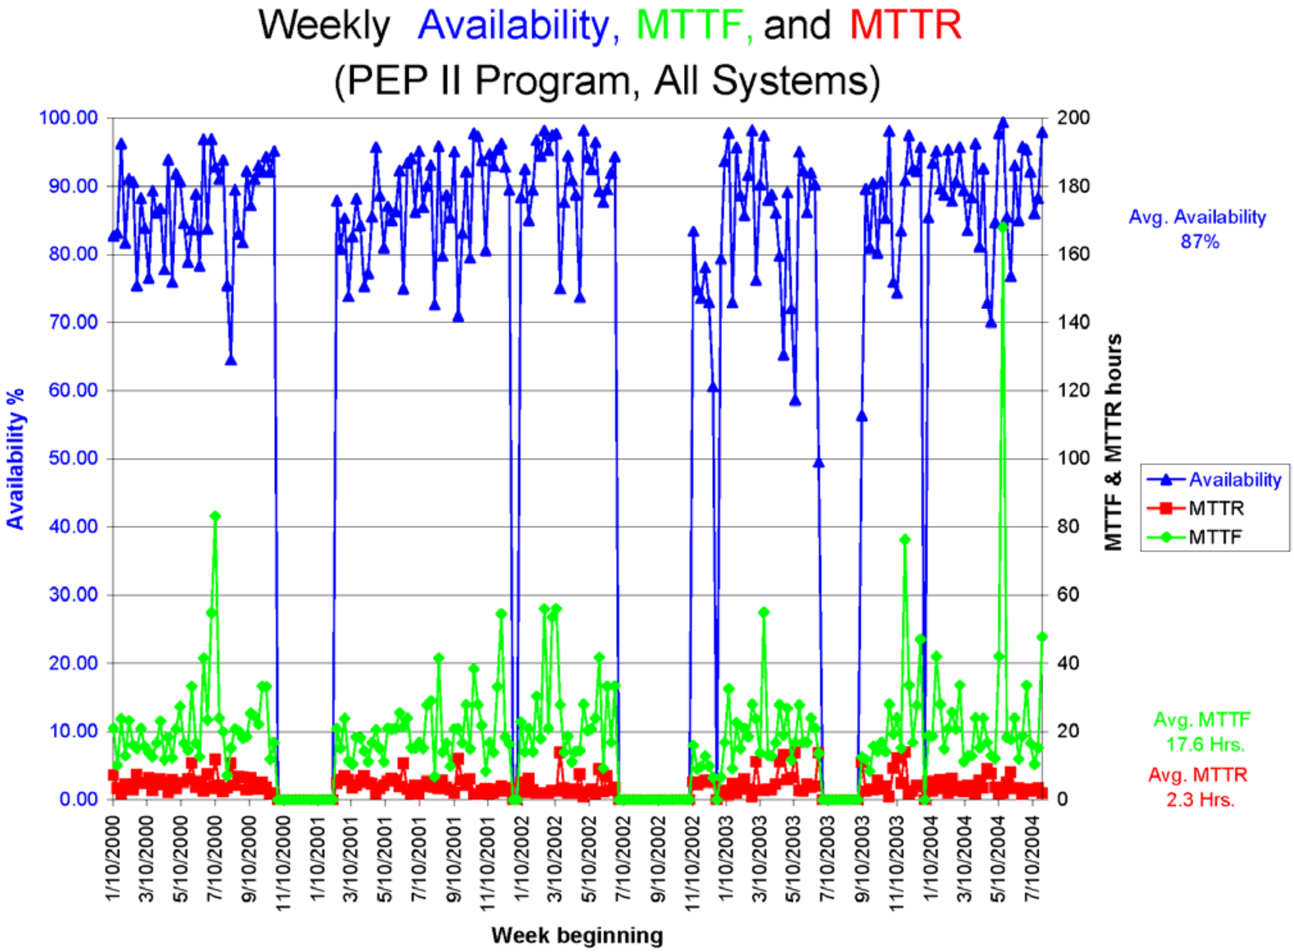

Fig. 1. Hardware availability, MTTF, and MTTR for the first four years of PEP-II production operations.

IV. Reliability by Subsystem

In Table II, the unscheduled down time is broken down by major system and by subsystem. Power supply problems were the worst offenders, accounting for 1010.8 hours of downtime, followed by RF systems, utilities, and controls. Other subsystems were relatively less troublesome. Among the RF systems, the longitudinal feedback systems accounted for nearly a quarter of the downtime. Note that the RF modulators are counted under power supplies. Among the power supplies subcategories, the pulsed power supplies were particularly troublesome, accounting for 148.2 hours of down time. It should be noted, however, that a significant fraction of this was due to initial difficulties with the new abort kicker systems at the beginning of Run 3. By the end of the run, these kicker systems were running reliably.

The ARTEMIS database enables the operations staff to evaluate the contribution of each type of system to the overall downtime. The database can be queried by any of its fields, such as by system, by area, by shop, or by date. Using this capability, staff members have observed interesting trends in the failure statistics of several subsystems. For example, the impact of vacuum system problems has changed dramatically since Run 2. In Run 2, 99.4 hours were lost due to 26 identified vacuum problems. In Run 3, 172 hours were lost due to only 6 problems. A cursory review of the specific problems revealed that most of the Run 2 vacuum problems were solved by brief entries to the PEP tunnel, with an average MTTR of 3.8 hours. A typical repair involved tightening flange bolts to stop a slow leak. By Run 3, however, the flange bolts had evidently all been tightened. Vacuum repairs were generally more difficult, sometimes requiring that vacuum chambers be replaced, resulting in an MTTR of 28.7 hours

TABLE II
DOWN TIME BY SUBSYSTEM

Subsystem	Subsystem Subgroup	Count	Downtime Hours
Power Supplies	Interlocks	70	161.7
	Controllers	43	76.6
	Power supply	293	498.4
	Cables	5	16.6
	Pulsed Power Supply	71	148.2
	PEP RF Modulators	72	100.5
	Sub-booster Modulator/Pwr. Sup	11	8.8
	Subtotal	**565**	**1010.8**
Magnets	DC	44	213.2
	Pulsed	4	12.6
	Magnet Mover	1	0.4
	Subtotal	**49**	**226.2**
RF	Linac Klystron	9	13.6
	Linac Modulator	27	30.9
	Sub-Booster	16	24.9
	Master Source-Drive Lines, PADs	3	9.3
	Sub Harmonic Buncher	3	8.6
	CW RF, DR RF & PEP RF	205	401.6
	Longitudinal Feedback	45	143.5
	Subtotal	**308**	**632.4**
Vacuum	Pumps	15	30.8
	Gauges & Controllers	12	13.8
	Valves	22	103.5
	Mechanical-Beam Pipes etc.	27	259.2
	Subtotal	**76**	**407.3**
Utilities	Electrical	61	308.7
	Water	119	237.4
	Compressed Air	6	27.4
	VVSs	1	0.7
	A/C & Chillers	11	43.1
	Subtotal	**198**	**617.3**

Subsystem	Subsystem Subgroup	Count	Downtime Hours
Cryogenics	Cryogenics Magnets	16	81.2
	Subtotal	**16**	**81.2**
		3	5.2
Controls	MPS (PLIC, PICs BIRs)	49	81.3
	PPS (BSOICs Keybanks, Doors & Gates, BTMs)	72	141.2
	Micros & Crates	95	115.8
	Networks	19	33.2
	BCS	5	7.8
	VAXs	8	11.4
	Cows	2	5
	Workstations	1	0.5
	Timing System	42	76.6
	BPMs	7	6.9
	Beam Monitors, Toros, Spectrum Foils	4	9.9
	Profile Monitors	1	1.5
	SWE (Software Installs/Changes)	37	46.9
	Vacuum I&C	30	34.5
	Subtotal	**375**	**577.7**
Safety	Fire Alarms	4	10.9
	Subtotal	**4**	**10.9**
		5	22.5
Other	Thermionic Gun	5	4.4
	Polarized Gun	12	38.7
	Experimental Equipment	72	146.9
	Equipment Checkout	11	44.8
	NTF	28	61.9
	Subtotal	**133**	**319.2**
	Total	**1724**	**3883.0**

TABLE III
EVENT COUNTS BY BEAM TIME LOST

Beam time lost	Events	Percent of total events	Hours down	Percent of down time
> 0 to 1.0 hours	859	49.8%	507.7	13.1%
> 1.0 to 2.0 hours	381	22.1%	616.3	15.9%
> 2.0 to 4.0 hours	287	16.6%	861.2	22.2%
> 4.0 to 8.0 hours	112	6.5%	632.4	16.3%
> 8.0 to 24.0 hours	74	4.3%	899.7	23.2%
> 24.0 hours	11	0.6%	365.7	9.4%
Totals	**1724**	**100.0%**	**3883.0**	**100.0%**

V. REPAIR TIME DISTRIBUTION

Table III sorts all events according to the beam time lost by each. Nearly half the events cause less than one hour of lost beam time, and together account for only 13.1 percent of the total lost time. 72 percent of the events cause less than 2 hours of unscheduled down time each. The remaining 28 percent of the events, causing more than two hours of down time each, account for 71 percent of all the unscheduled down time. Among these were 11 events that caused more than 24 hours of lost beam time each. These included five vacuum chamber failures, three site-wide electric power outages, two failures of a large dc power supply needed for a quadrupole near the collision point, and an SLTR quadrupole magnet coil which overheated and failed when a cooling water pump stopped. One of the power outages occurred when a tree branch grew too close to the 230 KV power line outside the SLAC perimeter fence. The PEP program lost 47 hours to the resulting power outage. Routine tree trimming had been deferred, but has now been reestablished.

VI. CONCLUSION

The statistics presented above are intended to provide a general overview of the hardware reliability of the PEP-II facility, and can be used as a basis for comparisons over time to look for trends. However, the specific numbers must be used only with some caution, because of unavoidable simplifications that have been done in these analyses. One simplification involves the counting of concurrent or overlapping problems. The statistics are based entirely on a database of malfunction reports generated mainly by the operations staff. In compiling the data for these analyses, we identify the time lost due to each malfunction. Sometimes, especially during start-up periods, machine operation can be obstructed by more than one problem simultaneously. In these cases, the down time is assigned to the first or most significant problem to avoid double-counting. Another simplification arises because some malfunctions allow the scheduled PEP program to continue at reduced luminosity while repairs are being done; hence many of the malfunctions in the injection systems do not show up as unscheduled down time, because the stored beams in PEP continue to circulate and collide while repairs are completed.

Nevertheless, the ARTEMIS database of hardware problems has proven to be valuable in identifying recurring problems and in quantifying the impact of various categories of problems. With this information, it has been possible to allocate engineering resources where they have had the most beneficial impact on the performance of the facility.

VII. ACKNOWLEDGMENT

We gratefully acknowledge the Operations staff members who collected the data used in this study and the staff members of the Stanford Linear Accelerator Center whose efforts in maintaining the linac and PEP-II facilities have provided the foundation for an extremely successful program.

VIII. REFERENCES

[1] C. W. Allen, S. Anderson, R. Erickson, W. Linebarger, J. C. Sheppard, and M. Stanek, "Opportunistic or Event-Driven Maintenance at the Stanford Linear Accelerator Center," SLAC-PUB-7424, March 1997.