Memory System Parameters:

L1 Cache - 0.5 ns access time - 95% hit rate
TLB - 1.0 ns to do a mapping - 90% hit rate
L2 Cache - 10 ns access time - 90% hit rate. (Both page table entries and actual data may be cached)
Memory - 100 ns access time - .01% fault rate for data accesses only (page table never faults)
Disk - 10 ms = 10,000,000 ns access time

L1	TLB	Page Tbl L2 Memory	Data L2 Memory		Probability	Time ✦
H					.95 = .95	.5 ns
M	H	(N/A)	H		.05*.90*.90 = .0405	1+10 = 11 ns
M	H	(N/A)	M	•	.05*.90*.10 = .0045	1+100 = 101 ns
M	M	H	H		.05*.10*.90*.90 = 00405	10+10 = 21 ns
M	M	H	M	Res	.05*.10*.90*.10*.9999 =.000449955	10+100 = 110 ns
M	M	H	M	Fault	.05*.10*.90*.10*.0001 =.000000045	10+10,000,000 = 10,000,010 ns
M	M	M Res	H		.05*.10*.10*.90 =.00045	100+10 = 110 ns
M	M	M Res	M	Res	.05*.10*.10*.10*.9999 =.000049995	100+100 = 200 ns
M	M	M Res	M	Fault	.05*.10*.10*.10*.0001 =.000000005	100+10,000,000 = 10,000,100 ns
					(Sum of the above = 1)	

• If a page's mapping is recorded in the TLB, it is must have been referenced recently and can therefore be assumed to still be resident, so assume faults don't occur in this case

✦Assume that when a miss occurs in a cache or TLB, the time is subsumed in time for the next level, and that when a page fault occurs the time to update memory when page is brought in is subsumed in the disk access time

AMAT = .95*.5 + .0405*11 + .0045*101 + .00405*21 + .000449955*110 + .000000045*10000010 + .00045*110 + .000049995*200 + .000000005*10000100 = 2.07 ns

(Revised 11/18/2019)