Example 2.4 A database server pulls 20,108 tpmC in a TPC-C benchmark. Each transaction induces 6 IOs at a disk which is capable of sustaining 240 IOs per second (IOps). How many disks should be configured for the benchmark? The system throughput in transactions per second (TPS) is:

$$X = \frac{20,108 \text{ tpmC}}{60 \text{ s}} = 335.13 \text{ TPS}$$

and using $V_d = 6$ IOs per transaction as the visit ratio in the forced flow law (2.7), the local disk throughput is:

$$X_d = V_d X = 6 \times 335.13 = 2010.8 \text{ IOps}$$

Since the maximum disk throughput is only 240 IOps, the minimum number of disks required is:

$$N_d^0 = \frac{2010.8}{240} = 8.38 \tag{1}$$

or 8–9 physical disks. But these disks would be maxed out at full utilization and could throttle the TPS rate. A more realistic sizing estimate follows from:

$$N_d = \left(\frac{Tx}{s}\right)\left(\frac{IOs}{Tx}\right)\left(\frac{IOps}{disk}\right)^{-1}\frac{1}{U_d} \tag{2}$$

where Tx/s the TPS rate, and U_d is the average utilization of each disk. To keep disk queues near zero length, disk utilization should be $U_d \ll 50\%$; which, as we shall see from (2.36), 50% corresponds to an average queue length of one. A more typical value is $U_d \simeq 15\%$. Substituting (1) into (2) produces:

$$N_d = \frac{N_d^0}{0.15} = 55.86$$

Therefore, the TPC-C platform requires at least 56 physical disks.