5. Conclusions and Future Research

In this brief chapter, we provide some thoughts on what has been and what remains to be accomplished in the areas we have studied. We organize our discussion by chapter.

Chapter 2: A model of an unreliable machine

We have made great progress for the case of i.i.d. exponential failures and i.i.d. exponential repair times, in deriving expressions for the amount of time required to produce a fixed number of parts, and for the number of parts produced over a fixed time interval. For both of these random variables, we have derived the moments, probability density functions, cumulative distribution functions and Laplace transforms (sometimes in terms of modified Bessel functions).

We have made very little progress when the assumptions on repair times and failure times do not hold. We note that in some cases, the repair time distribution may not be independent of the time since the last repair, and the time until the next failure may not be independent of the repair time. Further, although we believe (based on unpublished data) that the exponential distribution is reasonable for GM metal stamping lines, there will be situations in which a different distribution must be used. Barlow and Proschan (1965) and Proschan and Pyke (1967) describe a statistical method to test if data was generated by an exponential distribution.

Our models also assume that the parameters required are known exactly. It is easy to imagine situations in which the MTBF, MTTR, number of parts to be produced or the length of time available is not known with certainty. For example, the MTBF and MTTR might be uncertain if the machine, die, automation, jigs or fixtures, the part itself, or the operator are new or different. The number of parts that must be produced might be uncertain if some random quantity of parts that are produced are defective. The length of time available for production might be uncertain if, for example, the number of workers available is uncertain, or if the amount of machine time that must be reserved for other activities – such as preventative maintenance, changeover, or make-to-order parts – is uncertain.

Clearly, we have only scratched the surface. Many interesting and useful extensions to our results could be explored.

Chapter 3: Dynamic overtime decision model

We believe we have presented a useful operational model that could be used as part of a real-time decision support system to aid in the decision of when and how much overtime to run on an unreliable machine. We briefly comment on a few of the extensions that would make this model even more useful.

As mentioned in Chapter 3, the incorporation of stochastic demand would increase the number of production environments in which this model could be successfully applied. Our assumption that demand is known over some short horizon will not be true in some settings. We have made some progress toward this goal by studying special cases.

The addition of stochastic setup times to the model would better reflect reality in a metal stamping plant (i.e., see Chapter 4). Furthermore, the cost of overtime would ideally be a function of τ, since a setup crew may or may not be needed during the overtime shift, depending on the value of τ.

We might wish to model stockout costs as a function of the time that the order is outstanding, instead of incurring a one time penalty. This might be desirable simply because it is a better reflection of the cost structure incurred in a particular situation. However, we conjecture that the difficulties that we have with the optimal policy (the existence of a lower envelope, and possibly more than two critical values) would also be solved by such a change to the model. This would require new probability models in which we are able to compute the joint distribution of cumulative uptime and cumulative backlog.

Lastly, we note that an ideal model would be one that incorporates dynamic rescheduling. We expect this to be very difficult to achieve.

Chapter 4: Comparison of operating policies for a single unreliable machine

We will not repeat the conclusions that we presented in the final section of Chapter 4. We instead offer a few remarks about future research opportunities.

Our model assumed that demand was uncorrelated over time and across parts. Since this assumption does not hold true in many environments, it would be worthwhile to explore how, if at all, such correlations change our conclusions.

We note that P7 is a myopic policy, since it bases its production decision on the likelihood of stockout over the next production interval. A better policy might look ahead several intervals each time it must make a decision. An improved policy such as this might also be able to properly respond to cycles of varying lengths, which might reduce the undesirable behavior experienced at high utilizations.

Of course, we identified seven additional policies in Section 4.1 that may be worthy of exploration. Of particular interest is P8, which retains the desirable properties of P1 yet may perform better when demand variability is high. Policies of the type suggested by Graves (1980) that authorize production based on both aggregate and individual item inventory levels may also prove to be highly successful.

References for Chapter 5

Barlow, Richard E. and Frank Proschan. Mathematical Theory of Reliability. New York: John Wiley & Sons, Inc., 1965.

Graves, Stephen C. "The Multi-Product Production Cycling Problem". AIIE Transactions, 12, pp. 233-240, 1980.

Proschan, Frank and Ronald Pyke. "Tests for Monotone Failure Rate". Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume III, pp. 293-312. Berkeley, CA: University of California Press, 1967.