OPTIMAL PREVENTIVE MAINTENANCE SCHEDULING IN PROCESS PLANTS

LaRisa Sergent and Jeffrey Sorenson

Advanced Process Design CHE 4273

Dr. Miguel Bagajewicz

05/02/08

Abstract

This work expands upon a previous paper by Nguyen et al. which developed a new methodology for assessing the effectiveness of preventative maintenance scheduling guidelines on a chemical processing plant. Here, the model is applied to a larger plant than it was previously tested on and is expanded to include a risk analysis. Optimal parameters such as labor force size and preventative maintenance intervals are presented for the larger plant. The results of genetic algorithm testing are included to verify the results of the model optimization. Finally, plans for future work on this model are discussed.

I. Introduction

The cost of maintaining equipment in a chemical processing plant can range from 30% - 50% of the plant's total operating budget. With such a large amount of capital being placed in plant maintenance, it is imperative to establish an economically optimal maintenance policy. The goals of any policy are to maximize safety for both plant personnel and the public, and to minimize the total cost to the plant by decreasing loss of product and cost of maintenance. This paper will focus primarily on optimizing costs and determining how much preventative maintenance is economically optimal.

The cost of maintenance itself is significant, but when a piece of equipment fails, there are a number of losses which can result. Economic loss can be due to shutdown of the plant, replacement costs for the equipment, and decreased performance of failed equipment. Additionally, having the optimal number of maintenance workers and spare parts on hand can play a large part in cost optimization. Modern maintenance philosophy in chemical processing plants supports constant preventative maintenance (PM) rather than corrective maintenance (CM). However, due to limited time and resources, PM must be prioritized. Scheduling maintenance can be very complicated when pieces of equipment are interdependent. Effective maintenance scheduling is required to allocate resources and manpower in the most economically advantageous way.

All maintenance scheduling models use simplifying assumptions about the plant's equipment reliability, necessary downtime, product losses, and other variables. Mathematical models can be used to assess many decision variables simultaneously. Current optimization models differ somewhat in their assumptions and policies on maintenance, as well as in their ease of use. Genetic algorithms allow for the computation of more complicated models with few simplifying assumptions, and thus have been the latest trend in maintenance optimization. Monte Carlo simulations, on the other hand, can easily be used to estimate parameters within the greater model, but can optimize only a single variable at once. Both types of models will be used in this study.

II. Summary of previous work

We will be expanding upon a preventative maintenance model developed by Nguyen et al.[1] The model aims to address the question of how much maintenance is economically optimal, specifically at process plants which will have constraints on labor and spare parts. The model uses a Monte Carlo simulation to manage resource constraints, while assuming a standard periodic PM policy, as it is commonly used in the industry. The main purpose of Nguyen's study was to show "the effect of various decision variables on maintenance performance," rather than provide the complete optimization.

The primary decision variables which were incorporated into this model are:

- Labor availability – man hours are assigned to each CM repair based on the priority of the failure and to each PM case based on the PM interval. When the number of man hours available reaches zero, no more maintenance takes place.
- Failure modes – failures are sampled based on the probability that they were caused by a certain failure mode. The effects of the differences in failure modes are incorporated into the model outcome.
- PM frequency – the preventative maintenance interval which is to be optimized is expressed as a fraction times the MTBF (mean time between failures). PM should be frequent enough to prevent failures, but not so frequent as to incur high maintenance costs.
- Spare parts inventory policy – the trade-offs associated with storing parts for corrective maintenance are assessed by the model. Minimal spare parts are always available for PM; therefore, PM is constrained only by labor availability.

Using data from the Tennessee Eastman plant, Nguyen et al. investigated the effect that changes to the aforementioned decision variables had on maintenance cost. The "objective function" for the model was the total cost of maintenance, labor, parts, and economic loss in the specified time horizon. It was found that for the 19 pieces of equipment at the Tennessee Eastman plant, the optimal labor force is 3 maintenance workers, the optimal PM interval is 1 time the MTBF for each piece of equipment, and it is optimal to keep some spare parts on hand for corrective maintenance. The average objective value (total cost) for a two-year time horizon with no preventative maintenance was 1.66 million dollars.

III. Expanded results for the Tennessee Eastman plant

a. Changes made to previous model

A few changes were made to the previous study before it was expanded. First, it was observed that with time horizon set for only two years, most of the equipment was not receiving preventative maintenance even once during that time interval. Table 1 shows the mean time between failures values for each type of equipment in the Tennessee Eastman plant. Only the pumps and compressors, each with a mean time between failures of 381 days, will reach their MTBF during the two year time horizon. For this reason, the time horizon is set to ten years for all of the runs in the current study.

Equipment Type	MTBF (days)
Valve	1000
Compressor	381
Pump	381
Heat Exchanger	1193
Flash Drum	2208
Stripper	2582
Reactor	1660
Heat Exchanger	1193

Table 1: Equipment MTBF

Additionally, this study attempts to adjust the preventative maintenance interval (PMI) separately for each type of equipment. In the previous model, this interval was determined by a variable called an x-factor. An x-factor of .1 would mean that preventative maintenance was performed at an interval of $1/10^{th}$ of each piece of equipment's MTBF, while an x-factor of 2 would set the PMI for twice that of the MTBF. Using the same x-factor for every piece of equipment presents a couple of problems. Table 2 shows key economic data for a compressor and a pump, which both have a MTBF of 381 days. A compressor which fails at 381 days would incur a cost of almost $50,000 (CM cost + loss due to failure), while performing preventative maintenance once in that time period (xf = 1) would cost $61,000 (PM cost + production loss). A failed pump, on the other hand, would cost $6,410 to be repaired after failure and only $110 to maintain at a PMI of 381 days. Clearly, it would be more economically efficient to have a higher x-factor (i.e. more frequent maintenance) for pumps than for compressors.

	MTBF (days)	**CM cost ($)**	**PM cost ($)**	**Economic loss due to failure ($/day)**	**Production loss due to maintenance ($/day)**
Compressor	381	37,400	1,000	12,000	60,000
Pump	381	6,400	100	10	10

Table 2: A comparison of maintenance costs for two pieces of equipment with the same MTBF.

b. Optimizing x-factors by equipment type

As the economic comparison between compressors and pumps suggests, a single x-factor for all pieces of equipment is not optimal. Differences in maintenance costs, economic losses, and maintenance down-times indicate that what is ideal for one component may not be as ideal for another. To briefly study this issue, each component was individually optimized by plotting cost probability as a function of PM interval, setting the interval of all other components at the MTBF of the respective component and adjusting the chosen component from 0.4 to 1.4 times MTBF in intervals of 0.2. Figure 1, below, shows the effects on cost probability of changing the PM interval of the reactor.

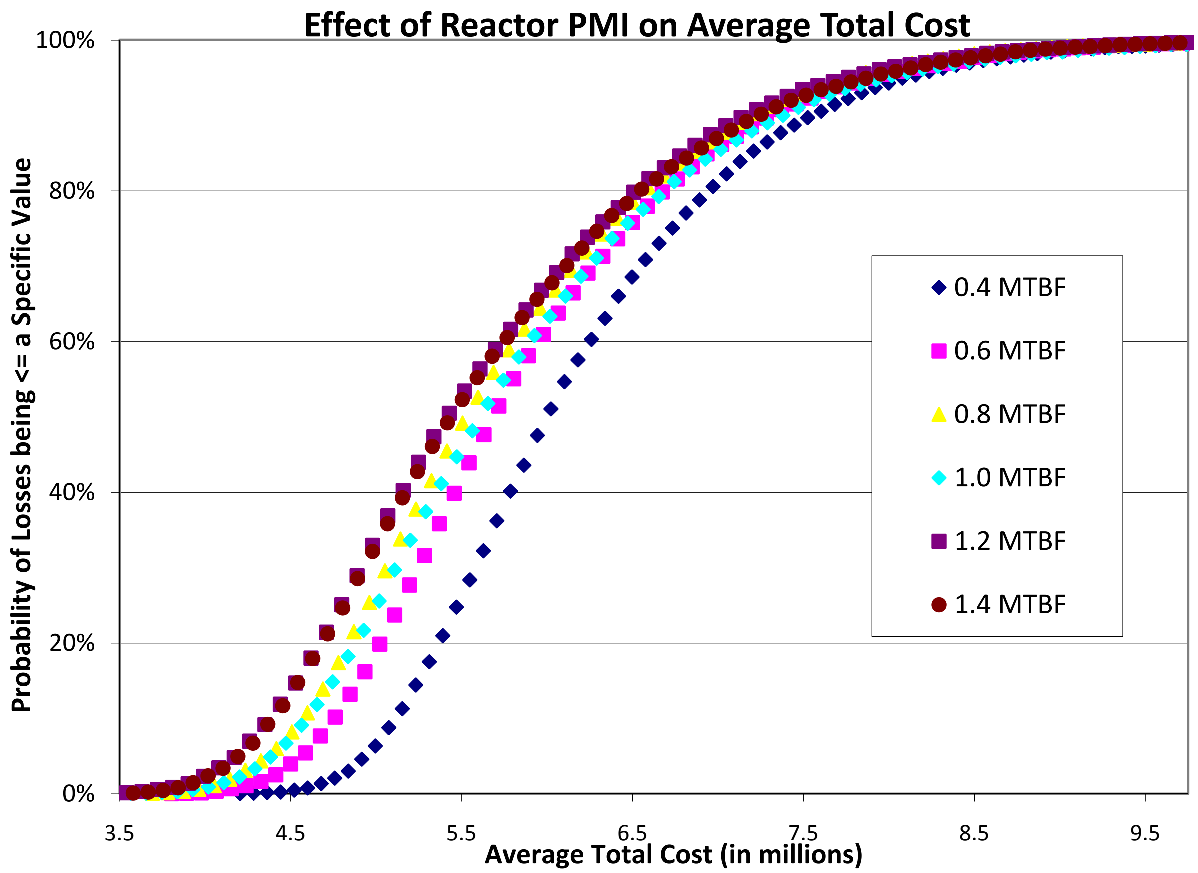

Figure 1: Optimizing by component type – reactor

As shown on this graph, a PM interval of 1.2 times the MTBF for the reactor is optimal for the process when all other intervals are set to 1.0. Repeating this process for all components, results were obtained as shown in Table 3. A model was then run to evaluate the cost probability using the optimized values for each component. This optimized model had a total cost of $5.66 million, just under the $5.67 million found optimal for a common x-factor of 1.5. However, optimizing by component was not evaluated for PM intervals beyond 1.4. It is highly probable that further evaluation of components at these higher intervals may further enhance the benefit of optimizing PM interval by component type. Utilizing a genetic algorithm to optimize all components simultaneously would also be of benefit. The theory and application of genetic algorithms are discussed later in this paper.

Component	Optimal PM Interval
Valve	0.5
Compressor	0.8
Pump	0.8 and 1.4
Heat Exchanger	1.0
Reactor	1.2
Stripper	1.0
Flash Drum	1.0

Table 3: Optimal PM Interval, by component.

c. Risk analysis

The Tennessee Eastman model was also expanded to include risk analysis. The objective value which the model determines based on the labor force size, maintenance intervals, and other variables which it is given is, of course, an average of many Monte Carlo simulation runs. Ten thousand samplings are averaged in this study to determine the objective value for any given set of parameters. In some cases, the values may range greatly, showing that a very large economic cost and very small economic cost are both possible. This situation would be one of great risk.

Figure 2 shows the risk analysis for various values of xf (in this case, there is a single x-factor for all pieces of equipment). The probability of occurrence on the y-axis corresponds to a value of total cost on the x-axis and includes the probability that the cost will be at that value or lower. Curves further to the left in the figure have lower average total costs. Curves that converge to 1 at a lower total cost value have less risk (lower probability) of high total costs. So, the optimal curve would be both steep and to the left. Here, an x-factor of 1.5 is shown to be optimum.

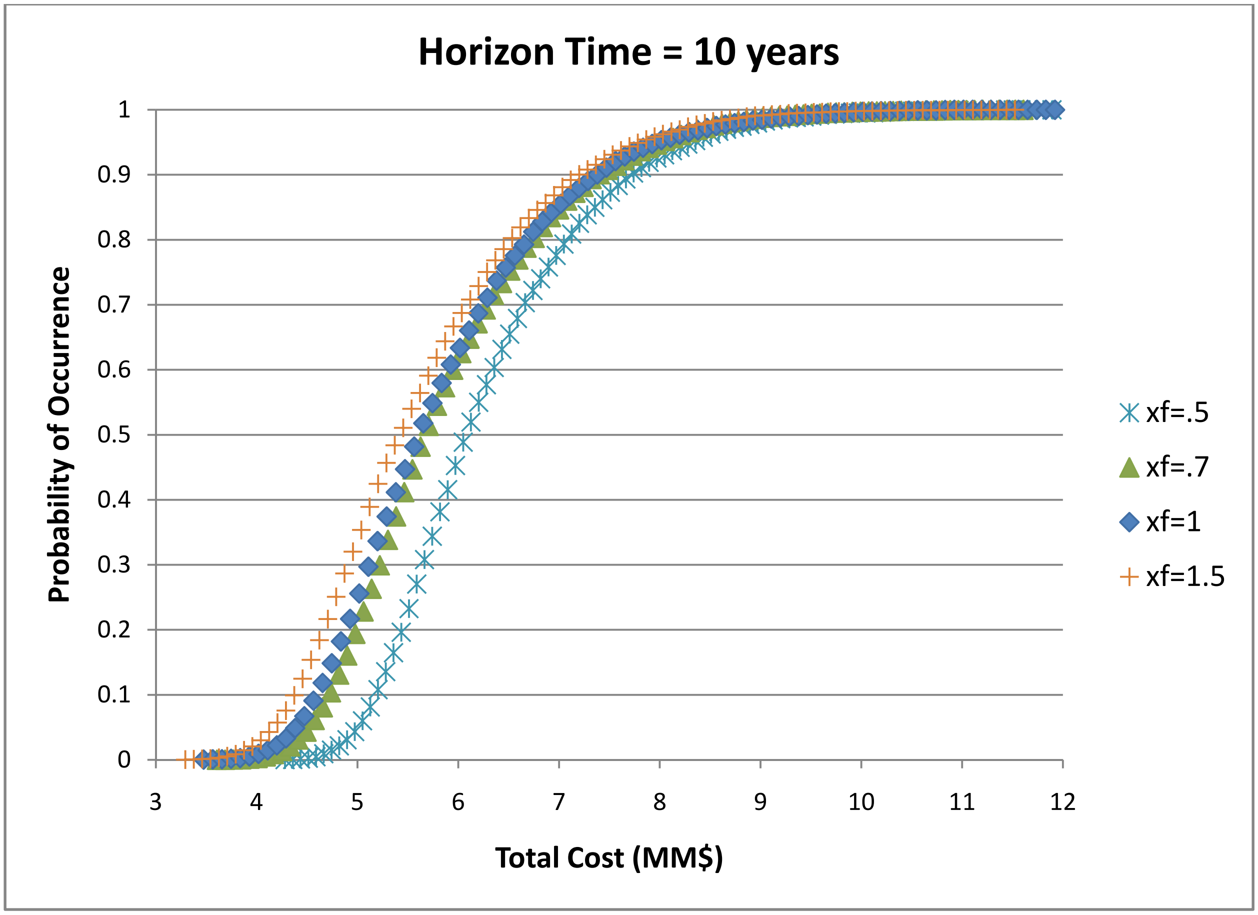

Figure 2: Risk analysis at Tennessee Eastman, based on varying x-factors.

In most cases involving risk, there was expected to be a trade-off between obtaining a lower average total cost and achieving minimal risk (as demonstrated in Figure 3). The red line represents a maintenance philosophy which would significantly favor preventative maintenance to decrease the possibility of dramatic losses in product. The downside to this approach is a higher average total cost.

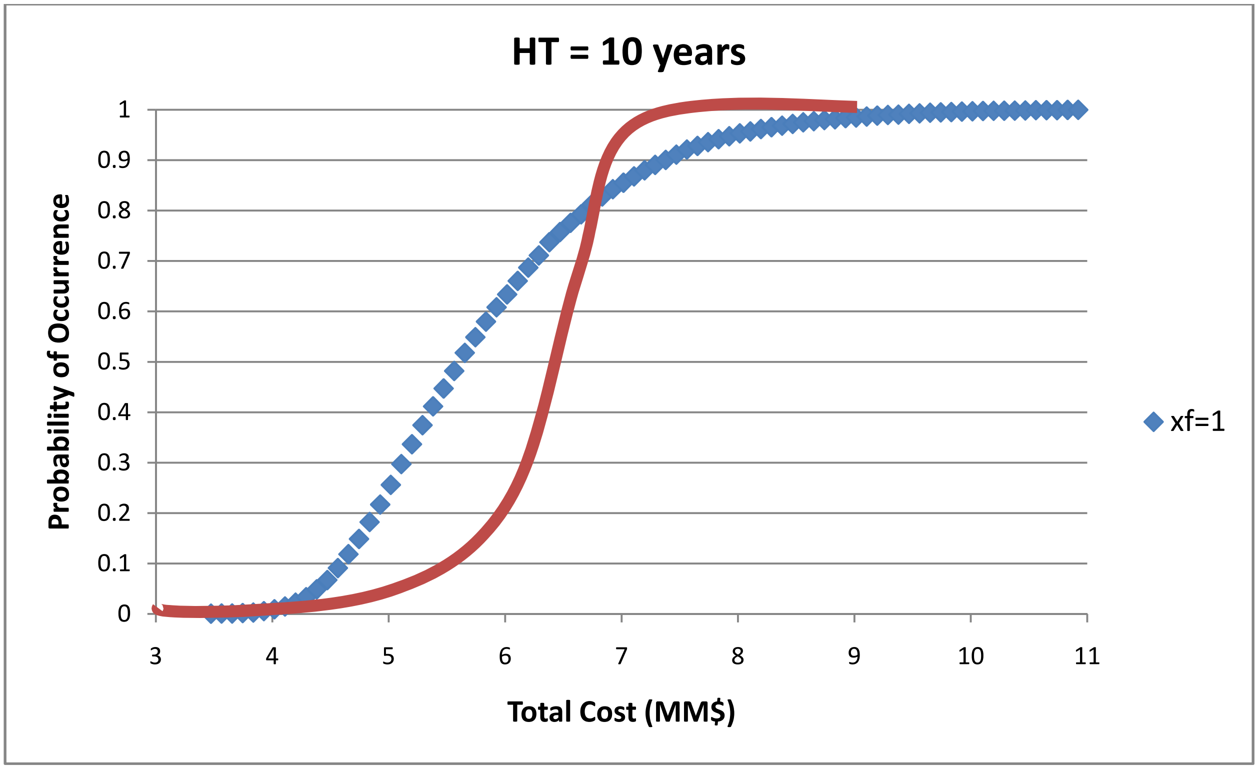

Figure 3: Expected trade-off between low ATC, high risk (xf = 1) and high ATC, low risk (some other xf).

This trade-off can be seen in the data for the Tennessee Eastman Plant, but it is very small. As shown in Figure 4, the curves for an xf of .9 and an xf of 1 actually cross each other. This means that although using an x-factor of 1 results in a curve further to the left (lower total cost), it does not reach a probability of 1 as quickly as using an x-factor of .9 does (so it's more risky). Again, the difference in this example is small, but we will also look at risk in a larger plant with more equipment interactions and scheduling parameters (Section IV).

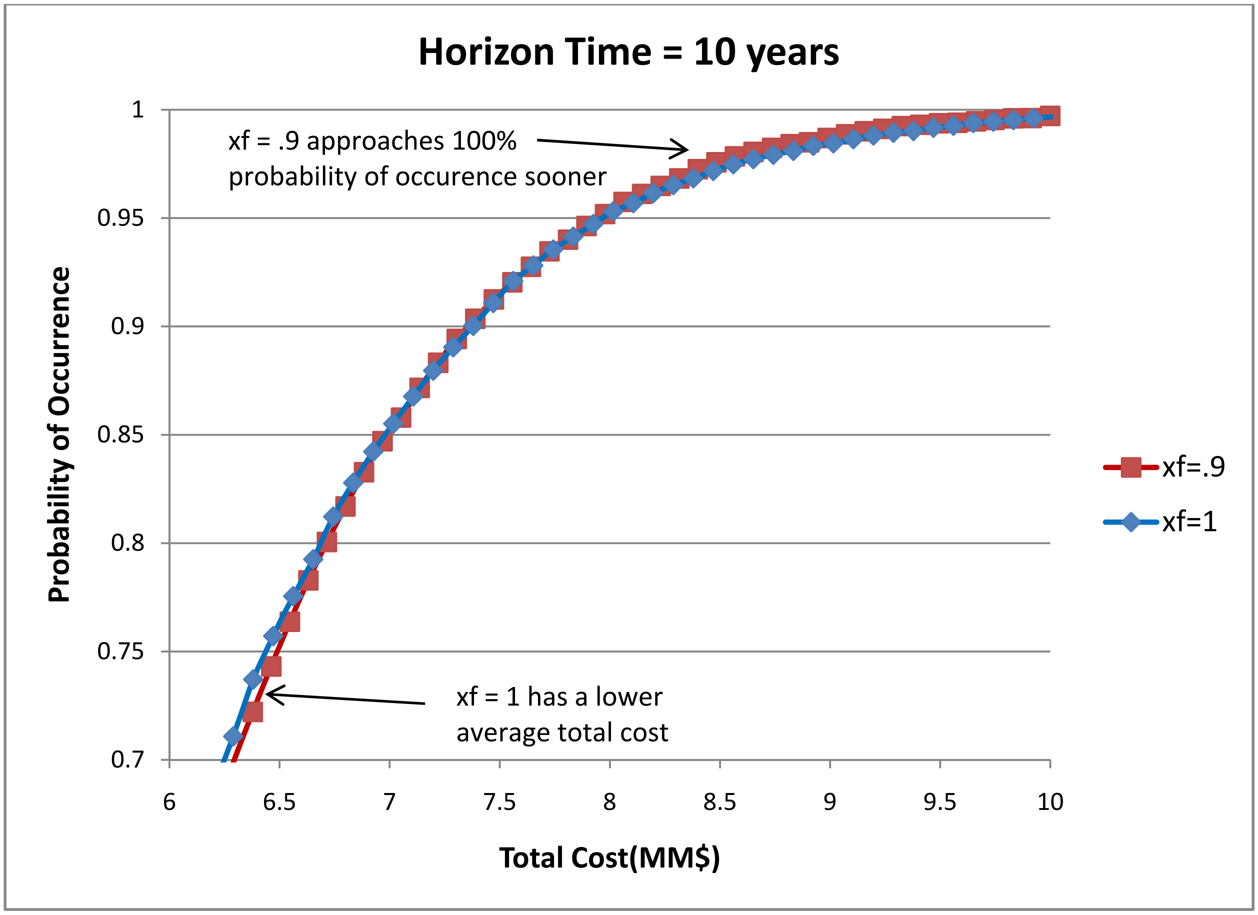

Figure 4: Risk and total cost trade-off between xf = .9 and xf = 1.

Because most of the risk analysis curves that were generated for the ten-year time horizon at Tennessee Eastman reached a probability of 95% at or around 8 million dollars, we chose this objective value as a benchmark for including risk analysis when determining an optimal preventative maintenance interval. The following figure shows both average total cost and percent of object values above 8 million as functions of the x-factor for PM frequency. Both of these functions can be minimized, but not at the same PM interval. The optimal PM interval for reducing both ATC and risk simultaneously is about 1.4 times the MTBF for this plant.

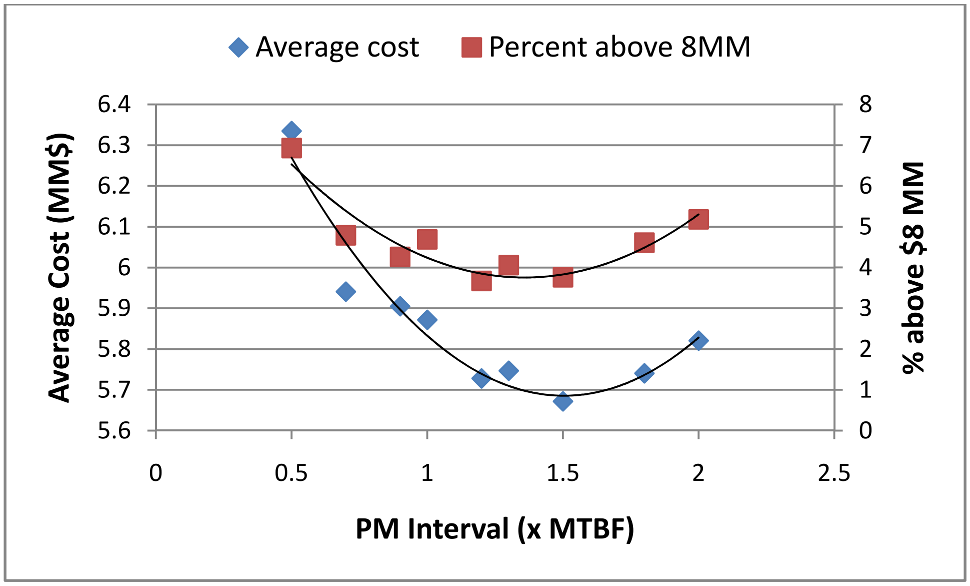

Figure 5: Average cost and risk of loss greater than $8 million for varying PMI at Tennessee Eastman.

IV. Large plant description

In seeking to expand the application of the model used in the previous study, it was necessary to obtain a larger process for the model to evaluate. A large west coast refinery volunteered equipment and volume specifications for its fluid catalytic cracker (FCC) unit. This unit, which processes roughly 50,000 barrels a day (bbl/day) of feed, allowed us to set-up the model to process almost 200 pieces of equipment, including 61 pumps (31 primary, 30 spare), 2 compressors, 4 heaters, 87 heat exchangers, 15 vessels, 1 catalytic reactor and its associated catalyst regenerator, and 12 columns and strippers.

MTBF for the equipment in the larger model has been assumed to be the same as that used in the previous study for similar component types (see Table 1). Heaters in the FCC unit were modeled as heat exchangers, while vessels were modeled based on flash drum data from the previous model. Due to their innumerability, valves were not considered in the FCC modeling.

Assumptions Applied in the FCC Analysis:

While additional information and research are needed to further refine the applied model, assumptions were made to determine values for economic loss and to determine what equipment needs to be placed into the maintenance model to keep it as simple, yet as complete, as possible. To this end, the following assumptions were made:

- Economic loss of product is assumed to be $10/bbl. This results in an economic loss of $500,000/day if the process unit is fully shut down.
- For the pumps in the process:

1. Spare pumps always work (no probability of failure).
2. If a spared pump fails, the spare instantaneously comes online (economic loss = 0).
3. After performing maintenance on a primary pump, the spare always undergoes PM (returns to good as new).
4. If a spare is insufficient to maintain a stream at its normal operating rate, economic loss is proportional to the loss in throughput.

These assumptions are reasonable and allow the model to be simplified by excluding the spare pumps from the equipment list used by the model. Their presence instead appears in the model by setting economic loss for pumps with spares as 0 (except as defined by assumption 4) and by increasing the assigned CM and PM costs to include the costs and time of performing maintenance on the spares. This reduces the number of equipment pieces the model must simulate for failure, increasing the program's efficiency without sacrificing function.

- For the heat exchangers and heaters in the process:

1. Failed exchangers transfer heat, but at a reduced rate (20-30% heat transfer loss).
2. Any exchanger located in series with other exchangers may be bypassed while being serviced without interrupting the process.

3. Economic loss is proportional to the portion of heat-duty lost due to the failure.

Further information is needed to know if exchangers in series can truly be bypassed for maintenance. Beyond this, the other assumptions are reasonable in that the same conditions may exist after an exchanger series with reduced area if the quantity of the stream being passed through is reduced to allow longer residence time for heat transfer.

- For the compressors, reactor, and regenerator:

1. If the component fails, the process goes offline.
2. The result is a maximum economic loss per day ($500,000/day).

This is reasonable, as losing any of these components prevents the catalytic reaction the reaction was established to perform. Loss of the regenerator prevents catalyst from being restored to further facilitate reaction of the feed, and loss of the compressor providing air to the regenerator prevents the catalyst from being reactivated by burning off the coke formed upon the catalyst surface.

- For vessels and columns:

1. Vessel failure results in failure of the associated process and, in most cases, shut-down of the unit.
2. Column failure results in shut-down of the unit.

These assumptions are reasonable in that loss of one step of the process, such as the depentanizer, would disrupt the ability to obtain products of the desired specifications from the end of the unit. The exception to this assumption is the failure of one of the "main columns" or their associated vessels. The feed from the catalytic reactor is split between two main columns, each producing identical product streams which are then rejoined and fed further forward in the process. Taking this redundancy into account, the assumptions for the main columns are:

3. When a main column fails or has a reduction in throughput due to failure of one of its associated components (vessels, exchangers, pumps, and so forth), the other column increases throughput as needed toward its rated maximum value.

4. The increase in throughput in a column will not extend beyond that needed to maintain the normal overall unit throughput.

This assumption reduces economic loss to the difference between the normal throughput and the sum of the new throughput values of each unit -- the functional one at (up to) maximum capacity, and the system in need of service at its reduced rate.

Again, much remains to improve upon these assumptions, including seeking feedback from the refinery as to their real-world representation.

V. Results of the FCC Model

After placing the information for the new equipment into the model, Monte-Carlo simulations were performed, as before, to determine optimal labor and PM intervals. With the program evaluating this larger process, the time required to perform each sampling increases to approximately 14 minutes, compared to the 2 minutes required to evaluate the earlier process. Therefore, an optimal quantity of labor and an optimal overall PM interval were determined rather than optimizing PM intervals for each type of equipment as done before.

a. Optimal labor force

PM interval was set to the mean time between failures for all equipment, while the units of labor were adjusted. Labor as a function of average total cost was found to be optimal with 5 maintenance employees. This is illustrated in Figure 6, below. As shown in the figure, adding additional units of labor increases the average total cost, though initially the value of economic loss continues to decline. While reducing economic loss is of value, this reduction in losses is insufficient to justify the costs of additional labor.

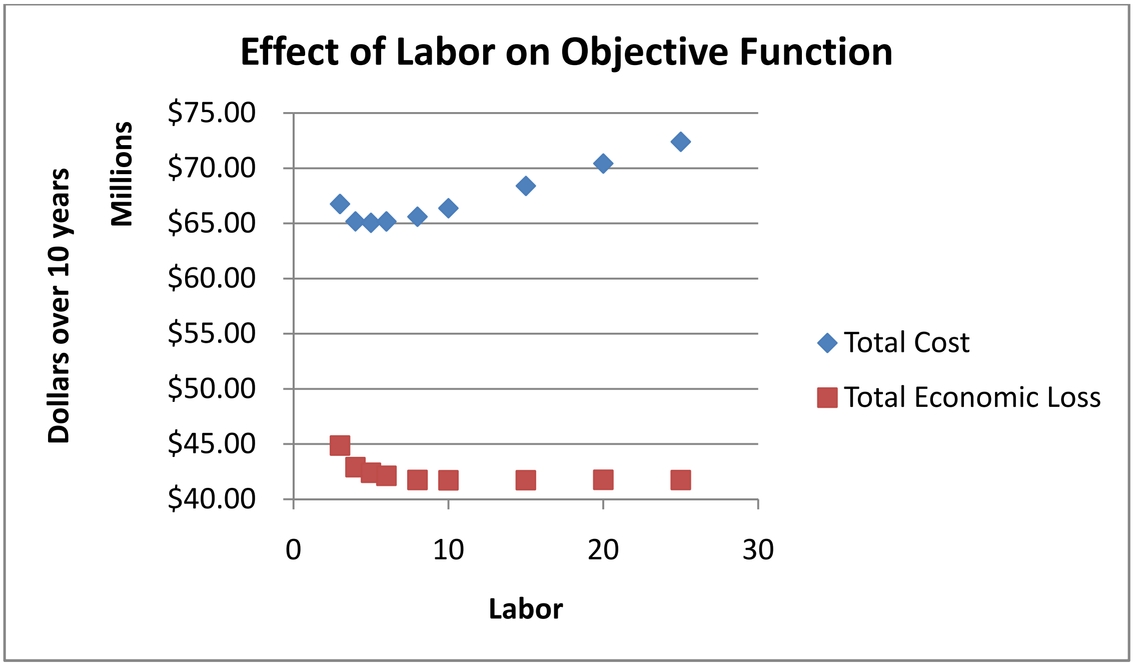

Figure 6: Total cost and economic loss as a function of labor for FCC

Overall, total cost is reduced by adding additional labor, as Figure 6 shows. Five maintenance workers are sufficient to handle both emergency CM and scheduled PM on the 200-some pieces of equipment at this FCC.

b. Defining a fitness function to evaluate risk

The effect of the PM interval was studied by varying the interval from 0.2 to 2.2 times the mean time between failures and plotting the associated average total cost (ATC), as well as by utilizing a fitness function to evaluate risk. This fitness function incorporates the objective of risk reduction by taking ATC and shifting the value by a fraction of the "Value at Risk," or VAR, for the process. Value at Risk is defined as the difference between ATC and the value where ATC has only a 5% probability of being higher. The fitness function, in turn, has the equation:

$$f_{fitness} = ATC + \alpha(VAR), \text{ where } 0 \leq \alpha \leq 1$$

The constraints upon α are due to the nature of the fitness function. If one desired to optimize based solely on average total cost, α would be set to zero. Similarly, if the only concern in optimizing is reducing the risk of high cost scenarios, α would be set at 1. For our analysis, a value of α=0.5 was used.

c. Fitness function and optimal PMI

The results of the comparison of the fitness function to PM interval are shown below in Figure 7. As can be seen, a minimum was found for a general PM interval of 1.7*MTBF, having an ATC of $64.97 million. While the function begins dropping noticeably beyond a PM interval of two times MTBF, this is due to the maintenance frequency being too long to evaluate pieces of equipment with high reliability (high numbers for MTBF) within the time horizon being evaluated. In order to obtain reliable values beyond this point would require expanding the time horizon further, at the cost of further increasing the time required to run a Monte Carlo simulation for each defined set of variables.

The similarity between the shape of the plot of ATC v. PM interval and the fitness function v. PM interval is due to the small rate of change in VAR as PM interval changes. While the value at risk would fluctuate, this variation was measured in tens of thousands of dollars, and as a result has little impact on function measuring change in millions of dollars.

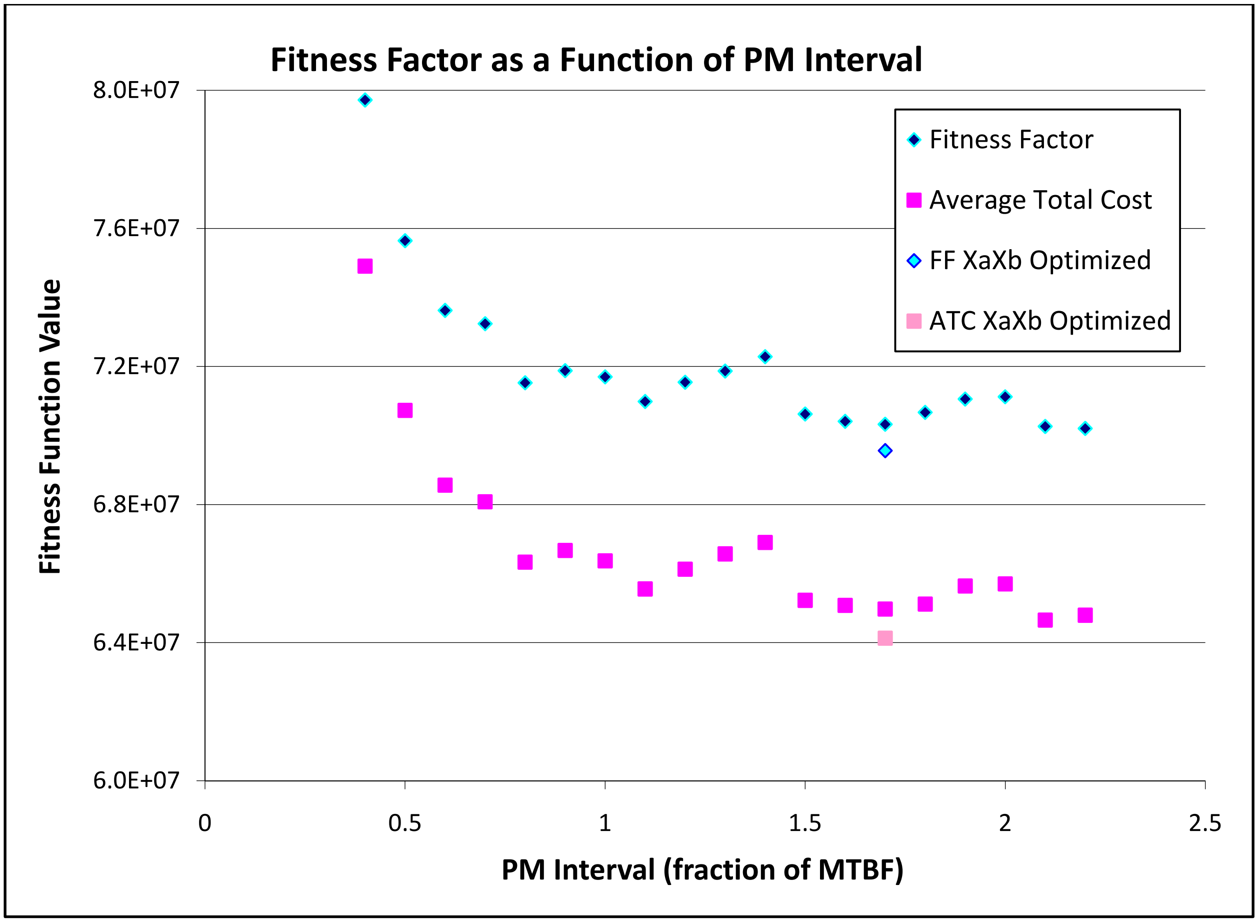

Figure 7: Average total cost and fitness function based on PMI for FCC

A lack of alteration in VAR is also apparent in Figure 8, where it can be shown that the cost probability function sees a horizontal shift at varied PM intervals, not a change in the shape of cost distribution. This is very similar to what was seen in the risk analysis plots for Tennessee Eastman. It appears that the most costly scenarios (approximately the top 5% of total cost probabilities) are the same regardless of the PM intervals. While adjusting PMI shifts the curves right and left – adjusting the average total cost – more frequent maintenance cannot prevent the plant from experiencing the “freak accidents” that are most likely responsible for the highest economic losses. Ultimately, for both Tennessee Eastman and the FCC, analyzing risk is not of value, and so economic decisions will focus subsequently on average total cost only.

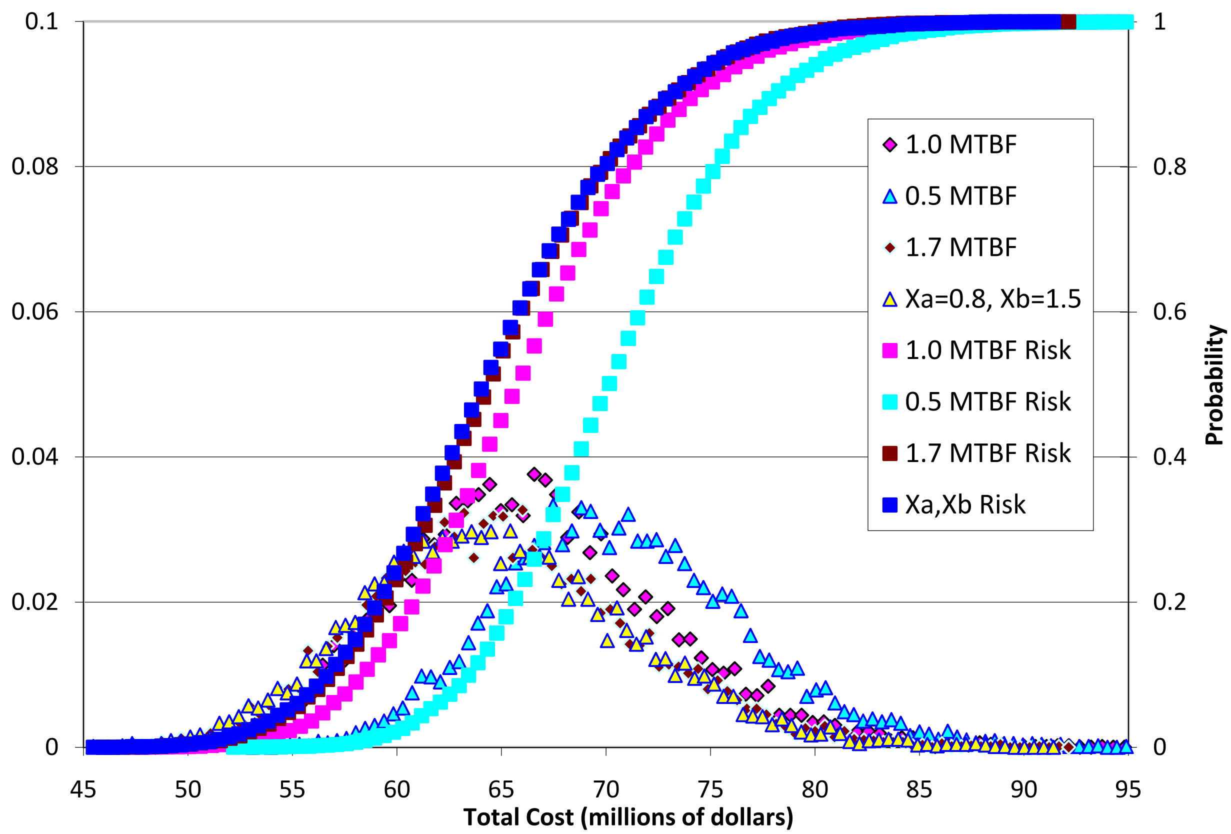

Figure 8: Cost-probability of FCC at varying PMI's

As a final point in the analysis, the optimal PM intervals for a two-equipment group system were evaluated for the FCC unit. For this analysis, the PM interval of the pumps was set to 0.8, and all other equipment was set at the "overall" optimal of 1.5 from the small-plant analysis. The results of this run are shown in Figure Y as the lighter colored point at PM interval = 1.5. This simple two-PM value model had an ATC of $64.13 million, $840,000 less than the best value found using a single PM interval. This value clearly indicates that, as before, costs may be lowered significantly by determining an optimal PM interval for each equipment group in the system. To assist in achieving this end, a genetic algorithm was applied to the model.

VI. Genetic Algorithm (GA) theory and results

a. Background and application

The Monte Carlo simulation model utilized in the prior results allowed total costs and risk of PM policy to be easily evaluated, but is limiting in that it requires the user to determine the PM interval for all groups of equipment in the process. While this is sufficient to establish an "over-all" PM interval for all equipment, or to evaluate the effects of changing PM policy for a single group while holding others constant, it is an undesirable and inefficient approach to obtaining a truly optimal solution. As hereto proven, various groups of equipment have varied ideal intervals for maintenance. However, it is impossible to assure the interval selected when adjusting only a solitary group will still be the most ideal if the intervals for any other group are changed.

To obtain this truly optimal solution using the manner previously discussed would require finding an initial "optimal" value for each equipment group, one at a time, by adjusting it while holding others constant. This process would then need to be repeated until it could be shown that no reduction in cost or risk could be established by changing any of the equipment group PM intervals. A much more reasonable and effective approach would be to use a program able to evaluate a model having a set of distinct PM intervals assigned for each group and able to compare that model directly to other models, each having a different sets of assigned PM intervals.

Toward this end, a genetic algorithm (GA) is used to study the same FCC process evaluated by the Monte Carlo model. The genetic algorithm set-up to evaluate this process is established by adapting the BASIC GA developed in the Bulgarian Academy of Sciences, Institute of Chemical Engineering and presented by Shopova et al.[2]

Genetic algorithms are named such for the semblance their functions have to biological processes. A brief summary of the algorithm operation follows:

1. The GA starts by creating a population of stochastically created models. Each variable in these models is assigned a random value from within a specified range defined that variable. We can regard the set of variables forming a model as a chromosome and the individual variables as genes on that chromosome.

2. The algorithm then evaluates each model/chromosome and assigns it a “fitness” value. In this case, the lower the average total cost, the higher the fitness value for each chromosome in the population.

3. The algorithm selects members of the population for reproduction, based on a random system weighted by fitness value. Chromosomes in the reproduction pool are assigned into pairs that serve as “parents.” Each pair of parent chromosomes produces two offspring. Offspring are exact copies of parents unless undergoing changes in the following two steps.

4. Produced offspring have a user-defined chance of “crossover,” where the values of one (or more) of the genes (a specific variable in the model) are swapped between the two offspring. Crossover, along with mutation, helps prevent the model from becoming trapped in a “local minimum.”

5. Produced offspring have a user-defined change called “mutation,” where the value of one (or more) genes is replaced by a newly determined random value.

6. The algorithm evaluates each offspring and assigns it a “fitness” value, as in step 2 for the original population.

7. The next generation population is selected from the pool of parent and offspring chromosomes. The chromosome with the highest fitness value is automatically placed in the next generation, and all other members of the new population are selected in a manner similar to step 3.

8. The algorithm repeats until it runs through the specified number of generations or the optimal value from each generation has not changed in forty generations.

At the conclusion of this process, the GA outputs the details of the found optimum solution, providing the PM interval used for each specified group of equipment in this optimal maintenance model.

b. Results, Part I

Initially, the genetic algorithm was set up with eight parameters – seven equipment groups and labor – and took 55-70 minutes per iteration (generation) to run. Constraints were placed upon the range of values for PM frequency of each maintenance parameter as follows:

Non-interfering pieces equipment: 0.05 – 1.1 x MTBF of equipment
Interfering pieces of equipment: 0.5 – 1.6 x MTBF of equipment

While genetic algorithms ideally require several hundred iterations to attain a solution regarded as highly optimized, time constraints at the time limited the the running of this algorithm to forty iterations. The GA was started separately on two different computers and ran over the following two days. The conditions for the algorithm are outlined in Table 4.

Population Size	40
Iterations	40
Crossover Probability	100%
Mutation Probability	30%

Table 4: Initial Genetic Algorithm Conditions

The best solution of each GA run are are presented in Table 5:

Variable Category	Run 1		Run 2	
	initial time	PM frequency	initial time	PM frequency
Group 1 (pumps)	1.2	0.45	0.7	0.8
Group 2 (compressors)	1.1	0.1	0.3	0.15
Group 3 (heaters)	0.7	0.85	1	1.3
Group 4 (exchangers)	0.3	1.5	1.4	1.15
Group 5 (vessels)	0.1	1.15	0.8	0.5
Group 6 (reactor/regenerator)	0.6	0.8	0.8	0.9
Group 7 (columns)	0.6	1.1	0.5	1.6
Labor	5		4	
Average Total Cost	$5,403,620/year		$5,590,003/year	

Table 5: Initial GA results for FCC plant

While both runs found solutions having a substantial cost reduction compared to the optimum conditions found during the Monte Carlo simulations, there is not a single common term between the two obtained solutions. Furthermore, while the average cost per year of these models is $5,496,812/year – a savings of over 15% from the value obtained for a single PM interval for all equipment – the lack of correlation between the two along with the substantial standard deviation of $132,000/year (2.4%) indicates that no truly optimal solution has yet been approached and further convergence of the algorithm is required.

Improving upon these results would require running the algorithm additional times, with a far larger number of iterations. Recognizing the difficulty of this task due to limited time, effort was taken to refine the algorithm to achieve convergence with fewer iterations and process each generation in less time.

c. Results, Part II

In working toward refining the genetic algorithm, equipment groups were reduced from seven to five by combining similarly modeled pieces of equipment into the same group. Heaters were combined with heat exchangers, and the reactor and generator were placed in the same group as columns due to the long MTBF time shared by both equipment types. We expected this reduction in the number of parameters to result in the algorithm converging to an optimal solution in fewer generations. In addition, the lower end of the constraint on non-interfering equipment PM interval was raised from 0.05 to 0.3 times the MTBF. This change in the constraint served two purposes: First, it removed the high PM frequency values that had shown themselves to be sub-optimal in the past, which would assist the model in converging faster. Second, the removal of these high-frequency PM interval solutions would allow the each model to be analyzed faster by the program, reducing the time needed to run each generation of the algorithm. Additionally, the cost of labor in the model was changed from $40,000 to $100,000 per unit per year, this new value being provided by the plant as the cost incurred for its labor force. The plant also confirmed that, excluding turn-around labor, the number of employees for FCC maintenance was indeed 5 – the same value previously found optimal in our Monte Carlo analysis.

Testing these changes over a few generations showed a time requirement of about 48-50 minutes per generation. The number of iterations was increased to 200, with the expectation that the genetic algorithm would be able to run this over approximately one week.

All other values for population size, crossover probability, and mutation probability remained unchanged. The results of the three runs of this modified genetic algorithm are presented in Table 6.

Variable Category	Run 1		Run 2		Run 3	
	initial time	PM frequency	initial time	PM frequency	initial time	PM frequency
Group 1 (pumps)	1.4	1.0	0.6	0.9	1.3	0.65
Group 2 (compressors)	0.8	0.45	0.2	0.4	0.2	0.4
Group 3 (heaters and heat exchangers)	1.0	1.4	1.4	0.9	0.9	1.15
Group 4 (vessels)	1.3	0.5	1.1	1.5	0.7	1.1
Group 5 (reactor, regenerator, & columns)	0.3	1.2	0.6	1.1	0	1.4
Labor	4		5		4	
Number of Iterations	96		62		105	
Average Total Cost	$5,767,811/year		$5,786,233/year		$5,756,541/year	

Table 6: Second-round GA results for FCC plant, including increased labor cost of $100,000/unit/year replacing previous value of $40,000

While set for 200 iterations, all runs "converged" much earlier than that. This is due to a condition built into the algorithm that if a more optimal solution is not found within 50 iterations of obtaining one, the program ceases running and proposes the found optimal value as a solution. As expected, the changes made to the algorithm resulted in greater convergence that previously obtained. With an average value of $5.77 million a year and a standard deviation of only $15,000/year (0.26%) across these three runs, it reinforces the fact that these are all solutions worth considering for real application, especially where common factors exist, such as in the case of the compressors. While only Group 2 has converged to a common value across all runs, the other equipment groups show far more similarity in their values than previously seen. Some of the remaining disagreement may be the result of combining multiple equipment types into single groups, in the case of Group 3, or an indication of a need to further classify an break-down an existing group, as in the case of Group 4.

VII. Plans for future work

This maintenance scheduling model is now able to provide a complete optimization of all the decision variables involved: size of labor force, preventative maintenance frequency, and spare parts inventory. The model has been applied to an even larger plant than was previously used, with a longer horizon time and more complex equipment interactions. Because we were expanding the model to work for such a large plant, a number of assumptions were made to more easily model the plant, specifically in calculating economic losses due to failed and being repaired equipment. In the future, it would be good to delve further into calculating more accurate economic loss values for each mode of equipment failure. Modeling the FCC plant in a simulation program such as PRO-II would allow for a better prediction of the interactions between equipment that would affect the total economic loss due to failure.

Additionally, future work could aim to program better assumptions into the model, specifically concerning shut down of the equipment. If taking a certain piece of equipment offline is going to cause the entire production for that day to be lost, then other pieces of equipment should receive preventative maintenance on that same day, rather than be taken down later and result in another total loss of production. It may be necessary to call in a "temporary labor force" in the event that many pieces of equipment need to be serviced in the same day. This would create a higher labor cost for the time horizon, but would likely result in a much lower objective function value, as less product would be lost due to the unavailability of being repaired equipment. As these principles are currently practiced in scheduling maintenance "turn-arounds" for large process units such as the FCC evaluated here, this would also assist the model to accurately reflect real-world scheduling practices.

Further progress may also be made with the genetic algorithm by expanding the number of equipment groups used and refining the constraints on the parameters of the model to more accurately represent the ideal ranges for each group. This would result in a greater correlation between the preventive maintenance intervals proposed in the solutions of each run, with the trade-off that such results would come at the inconvenience of a larger time-requirement for the genetic algorithm to run.

One final expansion which could be made is to incorporate safety as a specific model objective. Evaluating the tradeoffs between lower cost and higher safety would be difficult without assigning a dollar value to safety somehow. One possibility is to utilize a Pareto optimal curve, which is capable of comparing values with different units. A fitness function for safety could also be designed so that the weights placed on safety and cost could be determined on a case-by-case basis. Assessing safety itself would require data on plant incidents due to equipment failure, or some other way to categorize the hazards posed by a certain piece of equipment failing. This is by far the most difficult of our plans for future work.

References:

(1) Nguyen, D.Q., Brammer, C., Bagajewicz, M. *A New Tool for the Evaluation of the Scheduling of Preventive Maintenance for Chemical Process Plants*. University of Oklahoma, School of Chemical, Biological, and Materials Engineering. 2007.

(2) Shopova, E.G., Vaklieva-Bancheva, N.G. *BASIC – A genetic algorithm for engineering problems solution*. Computers and Chemical Engineering 30 (2006), pp.1293-1309