

Contents

Chapter 1

Introduction	1
---------------------------	---

Chapter 2

The Systems Engineering Approach	5
As Applied to Cogeneration Power Plants	5
Life-Cycle Cost and Modeling	7

Chapter 3

Considerations for a Cogeneration Facility	11
Conventional Power Distribution	11
Turnkey or Power Purchase Agreement	15
Time and Money	16
Territorial Behavior	17
Education	20

Chapter 4

Site Analysis	23
For Both New Projects and Existing Facilities	23
Data acquisition	23
Data analysis	23
Thermal Energy	36
Site Tour	37
Electrical switchgear	40
Noise	40
Dust	41
Working with the Local Utility and Interconnection Options	42
Supply and Demand	44
Generic Cogeneration Power Plant	46
Minimum Import Agreement	50
Inadvertent Export Agreement	51
Full Export, Net Metering, and Dispatch	52
Energy Storage	53
Island Mode, Intertied, and Island Mode Capable, or Off Grid Entirely	55

Island mode/off grid	56
Transients and starting surges	56
Continuity of power	59
Timing and balance	59
Reliability	61
Resynching to the utility	62
Fuel security	62
High latitudes and geomagnetic storms	63

Chapter 5

Power and Heat Sources	65
How Is It Done?	66
Reciprocating Internal Combustion Engines	67
Derating	68
Naturally aspirated or turbocharger-aftercooler?	69
Gearbox?	74
Cost drivers	77
Modeling	78
Overhauls	78
Oil changes	80
Reliability trees	80
Emissions levels	82
Rich or lean burn?	83
Engine and generator efficiency levels	85
Internal water pump?	86
Internal thermostats: How many, and how much flow do they control?	87
Costs	87
Combustion Turbines	88
Advantages of combustion turbines	88
Disadvantages of combustion turbines	90
Microturbines	94
Advantages of microturbines	94
Disadvantages of microturbines	96
Concentrating Solar Thermal	97
Parabolic trough thermal collector	97
Linear Fresnel solar thermal collector	99
Old-school technology with new bells and whistles	102
Thermal fluids, including organic Rankine cycle and molten salt	103
Costs	103
Locations and weather limits	103
O&M	104
Hybridization	104
Passive Solar Thermal—for Low-Grade Heat Recovery	104

Photovoltaic.....	105
Field sizing—area versus kilowatt output	105
Fixed versus single-axis or dual-axis tracking.....	105
Biomass	106
Biogas	107
Geothermal	109
Wind	110
Use where practical	110
Highly variable, therefore storage is needed.....	110
Other Sources	110
Hydroelectric.....	110
Fuel cell.....	111
Coal.....	112
Natural gas.....	112
Stirling engines.....	113
Energy Storage	114
Electrical energy storage.....	114
Mechanical energy storage.....	118
Thermal energy storage.....	119
N? Or N + 1?	120

Chapter 6

Distributed Generation Performance and Cost Analysis	121
Spreadsheet Models.....	121
Generic Model	122
Weather	124
Cloud cover	124
Randomness	126
Temperature	130
Microgrid Model	139
Financial Model	147
Reality versus Wishful Thinking.....	153

Chapter 7

Cogeneration Power Plants: Case Studies of Successes and Failures	157
Project 1.....	159
A rough start	162
Vortex versus differential pressure	163
Broken promises.....	164
First of a kind.....	165
Derating	165
Catalyst issues.....	166
Dampers.....	168

Coolant expansion tanks and dump radiators	168
Steam header	169
Air-conditioning	170
Operations-induced problems	170
Outcome	173
Project 2	173
Water where it's not supposed to be	174
Electrical distribution problems	177
Project 3	180
Electrical distribution problems	182
Steam piping (condenser)	184
Financial ramifications	187
Problems Common to Projects 2 and 3	188
Project 4	189
Application issues	190
Heat recovery	192
Containerized	193
Safety concerns	194
Project 5	195
Project 6	197
Duct burners	199
Operations and management problems	201
Project 7	201
Common Failures	204
1. Failure to evaluate all possible solutions to the situation	205
2. Failure to correctly identify the initial functions and requirements ...	206
3. Failure to address dust and other site-related hazards	207
4. Failure to size the generation units correctly or operate them	
in the correct mode	207
5. Failure to provide adequate engine operating environment	208
6. Failure to model all parameters that reduce engine operations	209
7. Failure to be realistic about operations and maintenance	210
8. Failure to consistently apply all rules and standards	
to each alternative	211
9. Failure to include life-cycle costs in financial analysis	212
10. Failure to perform or understand the results of sensitivity analysis ...	213
11. Failure to use realistic probability distributions in Monte Carlo	
analysis	214
12. Failure to document everything, including why a system was selected,	
for later reevaluation	215
13. Failure to identify the customer's true needs	215
14. Failure to perform failure analysis	216

Chapter 8

Concluding Points 219

- Benefits of On-Site Cogeneration 219
 - Reliability 219
 - Economics 220
 - Green 221
 - Easy, but not simple 221
- Build a Successful Team 223
 - Experience counts 223
 - Putting together a great operations and maintenance team 223
 - Training and qualification programs 225
- Excellence in Operations and Maintenance 228
 - Communication 229
 - Computerized maintenance management and preventive maintenance systems 229
 - Quality assurance and quality control 231
- Conclusion 232

Appendix A

Common Omissions and Failures 235

Appendix B

General Questions to Ask a Prospective Energy Services Company 237

Appendix C

True-False Questions to Ask a Prospective Energy Services Company 239

6

Distributed Generation Performance and Cost Analysis

Spreadsheet Models

Over the years, I have developed many spreadsheet models to try to figure out what my power plant designs would do over the course of their first year of operation (and beyond). With each new model, I added another new feature until I had included everything I could think of (and there's always room to add more), until my spreadsheets extended beyond column "DG" and were 35,065 lines long.

In this chapter, I will describe how I make my models. Because some of my calculations are proprietary, in these instances, I will tell you what I am doing and why, but you'll have to come up with your own formulas (sorry). Using data from existing power plants, I have validated these models by comparing their predictions to real outcomes.

Many factors govern how any given power plant will operate, and if these factors are not addressed individually, it can be difficult to analytically determine precisely why an existing plant is not operating as well as expected. For new power plants, or those that are designed but not yet built, failing to account for these factors will lead to guaranteed disappointment in the performance of the power plant (if it is ever built). Creating a detailed performance model will yield actionable results that will not only save you time and money but also provide a better product to your customer.

I use individual models for each different power and heat source, but they all include a provisions for a solar component, energy storage, and a prediction of the weather. Breaking down my model into specific areas of interest, I include sections for each 15-minute time interval covering:

- Time, date, and day
- Random numbers for variables
- A prediction of the weather
- An estimate of the net peak output of the generator(s), based on the weather

- The electricity imported by the prospective customer
- When to start an engine and load its generator
- Tracking the engine's load in specific ranges for maintenance purposes (low, partial, or full load)
- Using a protocol to unload the generator and shut down the engine if low load conditions are sustained until the desired threshold is reached
- Tracking the number of starts and stops of each engine
- Following the energy flows of the energy storage system, including the state of charge of the battery, number of charge/discharge cycles, among other parameters
- Calculating the new amount of power imported from the utility with the on-site power plant on line
- Determining the amount of fuel consumed by the engines
- Calculating the costs of the electricity for both the electric utility and the on-site power plant, as well as net savings to the potential customer

I then compare the results to the given on-site power generation problem. Ambient weather conditions end up playing a huge role in the performance of any combustion-based power plant, and the failure to account for weather-related effects at the beginning is the primary cause of disappointment at the end. Another variation to consider is using gensets in a peaker plant configuration as opposed to parallel operations. The benefits of each method will depend on the nature of the loads of the customer (and changes in those loads over the course of any given day). In any case, if you configure your model to allow for different operating modes, you will have additional information when comparing the various power plant configurations that meet your energy needs.

Given the size and complexity of the model described here, it will be helpful to create a spreadsheet and populate the cells with at least the column titles used with a note describing the purpose of those data. This can then serve as a template for future use.

Generic Model

The model I describe below is for a microgrid with three reciprocating engines and an optional solar (PV) field and an energy storage component (a battery). The first bits of the model are details I include in all my models, and the microgrid-specific bits will come into play later. I start the data-crunching parts, the meat and potatoes of the model, at line 15 or so. This allows room at the top of the sheet for defined data entry cells and summary information copied from the bottom of the sheet in a single place (e.g., annual engine run hours). Along line 15, I label the columns, and I leave line 16 blank so that I can insert preexisting conditions (e.g.,

whether the engines are running and the initial battery charge) before the first analytical line of data (which starts on line 17).

I usually start these models with three separate spreadsheets (arranged as tabs within a single file), between which data can be shared. These sheets are labeled “Model,” “Weather,” and “Comparison.” Since the weather profile affects everything else, I start with that, after the reference information that defines each 15-minute interval given in columns A–C.

Column A is the time reference, based on 15-minute intervals, since that is the interval the data from the potential customer’s electric utility is in (usually). I label cell A15 “Time” and A16 “0” (or leave it blank). Cell A17 is “15”; then A18 is “30,” A19 is “45,” and A20 is “1.” This continues with A21 being “15” through A24 being “2” and so on, all the way down column A to the end of the day, with cell A108 being “23” and A112 being “24.” A113 starts another day as “15” and the process continues down column A through the end of the calendar year with cell A35056 being “24.”

Column B is the date reference, with B15 labeled as “Date.” I leave B16 blank, as the “0” time reference in A16 is actually the end of December 31st (i.e., “31 DEC,” or whatever day/month designation you prefer) of the previous year. B17, at the 00:15 time stamp in column A, is the first actual data point in the New Year, which is “01 JAN.” B113 becomes “02 JAN,” aligning with the “15” in cell A113, and this is continued through to the end of the year.

Column C is the day reference, with C15 being “Day.” Having the day of the week as a reference is very handy, as it is helpful to differentiate weekdays from weekend days (as well as denote holidays). Most electric utilities have different rate structures for the different days of the week, and these will be incorporated into the economic section of the model, which is described later in this chapter (see the “Financial Model” section). Using a calendar for whatever year of data you’re analyzing, insert the appropriate day in cell C17, C113, and so forth, all the way down to the bottom of the sheet. These first columns are shown in figure 6–1. Also note that, although not shown in this example, I highlight the weekend and holiday data in bright color to make it easy to notice when scrolling through the data set.

14	A	B	C
15	Time	Date	Day
16			
17	15	1-Jan-15	Thur.
18	30		
19	45		
20	1		
21	15		

Fig. 6–1. “Time,” “Date,” and “Day” columns

Weather

Where you put the weather information in your model is not important—what is important is simply that you have a weather section, because the ambient conditions on any given day will have a significant impact on the performance of the power plant. If you are planning to enclose your power plant in a temperature-controlled environment, this bit of the model becomes very simple. However, if your power plant will be exposed to the ambient outdoor temperature, even if protected from the elements, the impact of the weather on the operation of the equipment must be taken into account. Most airports keep a log of the local weather data; conveniently, it's common to find these data in 15-minute intervals. Digging through the data on the [weather.gov](http://www.weather.gov) website can reveal a treasure trove of useful information.

As a bare minimum, include the daily high and low ambient temperatures in your model, because these factors have tremendous impacts on the performance of any power plant that uses the combustion of a fuel as the primary source of power. I also include cloud cover predictions to model the amount of solar radiation reaching the ground; these calculations may sound like voodoo, but bear with me as I describe the method behind my weather predictions. Thanks to the Internet, I can usually find the following information for the site I am evaluating (or local data close enough to the site), often as a graph of the annual data:

- Median cloud cover (percentage)
- Cloud cover types (clear, mostly clear, partly cloudy, mostly cloudy, overcast)
- Probability of precipitation (percentage)
- Types of precipitation (drizzle, light rain, moderate rain, heavy rain, thunderstorms, snow)
- Relative humidity (percentage)
- Wind speed
- Fraction of time spent with various wind directions
- Hourly temperature bands

Cloud cover

From the website weatherspark.com, I have adapted a graph of cloud cover types versus months of the year, given as a probability (percentage) for the location being analyzed. Within it, I divide each month into quarters and obtain the percentages for each type of cloud cover for each quarter-month period. Making any given month equal to four weeks allows easier interpretation of the graphical weather data. I put this information into a new spreadsheet (“Weather”), as a tab on the model, which will ultimately be only one of many separate spreadsheets that make up the performance model. This information is shown in table 6–1.

Table 6–1. Cloud cover probabilities

Month	Quarter	Overcast	Mostly Cloudy	Partly Cloudy	Mostly Clear	Clear
1	1	46	57	68	75	25
	2	45	56	67	75	25
	3	45	56	67	75	25
	4	43	55	66	75	25
2	1	40	53	65	74	26
	2	38	51	63	72	28
	3	34	48	61	71	29
	4	31	45	59	69	31
3	1	28	43	58	68	32
	2	27	42	56	67	33
	3	25	39	54	65	35
	4	24	37	53	64	36
4	1	21	36	49	61	39
	2	19	34	48	59	41
	3	17	30	45	57	43
	4	15	28	44	56	44
5	1	14	25	40	53	47
	2	13	24	38	52	48
	3	10	20	34	47	53
	4	8	17	30	44	56
6	1	7	15	27	40	60
	2	6	14	25	38	62
	3	5	12	23	36	64
	4	4	10	21	34	66
7	1	4	9	19	32	68
	2	3	8	19	30	70
	3	3	8	19	30	70
	4	3	7	18	30	70
8	1	3	7	18	29	71
	2	3	8	18	29	71
	3	4	8	19	29	71
	4	4	9	19	30	70
9	1	4	10	20	31	69
	2	4	10	21	31	69
	3	5	12	23	33	67
	4	6	13	25	34	66
10	1	7	16	27	37	63
	2	10	19	30	40	60
	3	14	24	36	46	54
	4	17	28	39	48	52

Month	Quarter	Overcast	Mostly Cloudy	Partly Cloudy	Mostly Clear	Clear
11	1	20	30	43	53	47
	2	23	35	47	56	44
	3	30	41	54	62	38
	4	32	44	56	64	36
12	1	38	49	60	68	32
	2	40	50	62	70	30
	3	44	54	66	73	27
	4	45	56	67	74	26

Estimating as best I can, I also generate a probability distribution of the amount of solar direct normal irradiance (DNI) that can penetrate the given cloud layer (as a percentage of the total possible amount) using a defined high and low limit. This is shown in table 6–2.

Table 6–2. Range of DNI penetration

	Overcast	Mostly Cloudy	Partly Cloudy	Mostly Clear	Clear
High	30	50	75	85	100
Low	15	20	40	60	85

Even though I have taken a few meteorology classes, I am certainly not a meteorologist; nevertheless, the values in table 6–2 are derived from personal observations, estimating the amount of sunshine that actually reaches the ground (or a solar [PV] panel) over the course of a year. Even on a “clear” day, there can be enough haze in the air to act as a filter, attenuating the DNI by up to 15%—or even more if you’re analyzing a site in (or near) a major metropolitan area, such as Beijing (for an extreme example). This is why the “Clear” column has a range of 85%–100%. Thus, although the values in table 6–2 represent my estimates of hourly averages of the DNI conditions on the ground for each sky coverage for central California, you should tailor your values to match your given location.

In any case, the values in table 6–2 reflect my estimates of the range of available sunlight reaching a PV panel or a parabolic trough given the rough category of cloud cover predicted from the random number generator. Yes, this method of prediction may seem like voodoo, but at the end of the year, the averages will have worked themselves out with offsetting errors, yielding a net result that should be fairly close to the real conditions (even those yet to be experienced).

Randomness

To keep things simple, I treat the weather at any given time as a random variable, but within limited ranges to cover an hourly average; there’s not much

point in trying to be more precise than that (similar to dealing with significant figures in engineering calculations). I don't think it's reasonable for the model to allow for a fully overcast sky in one 15-minute interval and a perfectly clear sky in the very next interval; however, an hour after initially overcast conditions, the weather could be completely different. If you've ever lived in Denver, you may have heard the expression, "If you don't like the weather in your front yard, go to your back yard—it'll be different by the time you get there." The basic idea is to use the spreadsheet's built-in random number generator to give you a value as a starting point, assigning a cloud cover amount to that value and then using the random number generator again to yield a figure that is between the two limits of net DNI available for that particular cloud cover assignment.

Looking at the first quarter of the first month of the year in table 6–1, we see that there is a 46% chance that it is overcast. The values in the table are cumulative. Thus, given the 46% chance of overcast conditions and an 11% chance that it's mostly cloudy, there is then a 57% chance that it's either mostly cloudy or overcast. Furthermore, you could include the 11% chance that it's partly cloudy and the 7% chance that it's mostly clear, to round things out; subtract that total from 100%, and what's left is the chance that there is a clear sky that day (25%). From this, we can allow a randomly generated number between 1 and 100 give us what the sky coverage is (based on what range the number falls into), and then a second randomly generated number can be used to determine the range of attenuation for each category of sky coverage to give the actual (predicted) percentage of the incident DNI reaching the ground or the PV panel.

Along line 15 in the model, starting with column A, we now have "Time," "Date," and "Day," so cell D15 can be "Random." From cell D17 (D16 is blank) down to the bottom of the data, the cells are "=RANDBETWEEN(1,100)." These values will change every time the return key is hit, which has its purpose, but it's also good to have a fixed set of randomly generated numbers for analysis, so I label cell E15 "Fixed." To populate this column, simply copy the values in the "Random" column (D) and copy into the "Fixed" column (E) as either values only or values and number formats. Cell F15 is "Wx" (weather), and for the values in this column we must reference the weather table (table 6–1) on the "Weather" tab, which occupies columns A–G on that sheet.

For example, for the first quarter of January (approximately seven days), I use the following formula for determining the sky cover in each cell in column F:

```
=IF(E17>Weather!$F$2,Model!$K$15,IF(Model!E17>Weather!$E$2,
Model!$J$15,IF(Model!E17>Weather!$D$2,Model!$I$15,IF(Model!E17>
Weather!$C$2,Model!$H$15,$G$15))))
```

In other words, I compare the randomly generated number in cell E17 to the cumulative percentages in the weather table, as shown in table 6–3.

Table 6–3. Cloud cover, the first week of January

Month	Quarter	Overcast	Mostly Cloudy	Partly Cloudy	Mostly Clear	Clear
1	1	46	57	68	75	25

If the randomly generated value is above 75, it is a clear day. If not, then is it above 68? If so, then it must be a number from 69 to 75, and it's mostly clear. If it is not above 68, is it above 57? If it is, then it's a partly cloudy day, and if it's not, then is it above 46? If it is, then it's from 47 to 57 and therefore mostly cloudy. Finally, if it is 46 or below, then it is an overcast day. By this method, every 15-minute interval will have its own randomly generated cloud cover, based on a set of limits for that particular week (or quarter-month), throughout the year.

To make things a bit more realistic and to smooth out some of the extremes of the randomness (e.g., overcast for 15 minutes, clear skies during the next interval, then back to overcast), I average the results over each hour. For more column headings, I label the cells G15–K15 as follows: “Overcast,” “Mostly Cloudy,” “Partly Cloudy,” “Mostly Clear,” and “Clear.” To estimate the solar DNI attenuation for each sky cover condition in cells G17–M17 (and continuing below, to the bottom of the sheet), I refer to the result in the “Weather” cell (E, Wx). The first several 15-minute intervals are in table 6–4.

Table 6–4. Cloud covers

Random	Fixed	Wx
27	51	Mostly cloudy
76	12	Overcast
69	100	Clear
87	96	Clear
22	11	Overcast
5	40	Overcast
10	32	Overcast
45	42	Overcast

In the space above line 15 in columns F–K, I create a small table with the ranges of DNI attenuation factors that are associated with each of the cloud cover possibilities. Just because it's a clear day doesn't mean there will be 100% optical transmission through the atmosphere. Haze, smoke from nearby wildfires (which happens a lot in California), and other factors can still reduce the DNI reaching your PV panels, even though it's technically a “clear” day. (This was presented in table 6–2.) In other words, if the random number generator's input to the cloud cover table shows that a particular 15-minute period will have

partly cloudy conditions, then the amount of DNI reaching the ground will be between 40% and 75% of the full amount possible.

The next step is to continue to populate the model with specific DNI attenuation factors that will be between the high and low limits established in table 6–2. There should be only one entry, with the other cells being a zero, as shown in table 6–5.

Table 6–5. DNI attenuation examples

Wx	Overcast	Mostly Cloudy	Partly Cloudy	Mostly Clear	Clear
Mostly cloudy	0	23	0	0	0
Overcast	18	0	0	0	0
Clear	0	0	0	0	93
Clear	0	0	0	0	97
Overcast	19	0	0	0	0
Overcast	22	0	0	0	0
Overcast	27	0	0	0	0
Overcast	19	0	0	0	0

In column L, I summarize the weather factor (denoted “Wx Factor”). Cell L15 is labeled with this column heading, and each of the cells below, to the bottom of the data in column L, is simply the sum of the cells in columns G–K. Since only one of those cells should be populated, this also makes for an easy cross-check that the model is behaving as desired. For example, cell L17 is:

$$=SUM(G17:K17)$$

So now we have an estimate of the random nature of the weather for each 15-minute interval for the entire year. Because the weather usually doesn’t change from fully overcast to absolutely clear skies as quickly as 15 minutes, I attempt to smooth the differences among the data into an hourly average. Over the course of a full year, this helps to smooth out the outlying data points into a more realistic representation of natural weather variation—and its impact on solar power sources. This is shown in column M (labeled “Wx Fx Avg.”), and it is simply the average value of each block of 4 interval data, with cell M20 being the first populated cell (table 6–6).

Table 6–6. Weather and DNI attenuation

D	E	F	G	H	I	J	K	L	M
Random	Fixed	Wx	Overcast	Mostly Cloudy	Partly Cloudy	Mostly Clear	Clear	Wx Factor	Wx Fx Avg.
15	51	Mostly cloudy	0	47	0	0	0	47	
37	12	Overcast	20	0	0	0	0	20	
72	100	Clear	0	0	0	0	100	100	
19	96	Clear	0	0	0	0	99	99	66.5
84	11	Overcast	28	0	0	0	0	28	
29	40	Overcast	21	0	0	0	0	21	
85	32	Overcast	21	0	0	0	0	21	
93	42	Overcast	24	0	0	0	0	24	23.5

I learned a long time ago that, even though the formulas are pretty simple and the odds of an error are low, it is practical to insert cross-checks into every spreadsheet model. There are many ways to do this, and the built-in formulas in most spreadsheets are helpful (especially when combined). In this case, since only one cell in columns G—K should have a number greater than zero, one way to verify results against expectations is to use a column labeled “Check” (e.g., column N), and have the cells in this column use the following formula (e.g., cell N17):

$$=COUNTIF(G17:K17,">0")$$

Since only one cell in the G—K range should be above 0, the result should be a 1 in each cell in the N column. At the bottom of the data set, another check can be made to verify that there are no cells in the column that have a number above 1 and another to verify no cells are below 1. When the results show anything unexpected, the use of conditional formatting of the cells in the area of concern (to highlight the errors) makes it easier to find them. Now, the forecast of actual DNI reaching the ground is complete. When some solar energy is considered in the hybrid mix of power and heat sources, a reasonable estimate of production can be made—and it will be a lot less than what the nameplate information predicts.

Temperature

The ambient temperature at any given moment has a tremendous impact on the performance of the equipment used in a cogeneration power plant. Consider the effect of air temperature on the heat rate of an engine or combustion turbine, as well as the follow-on effect of changing the costs associated with the fuel purchases; the effect of the difference (delta) between the temperature of the air and of the