Convex Capital Cost Optimization: A Rwandan Case Study

Phil Ebben
Department of Mechanical Engineering and Sloan School of Business
Massachusetts Institute of Technology
Cambridge, MA, USA
Philip.ebben@sloan.mit.edu

I. Introduction and Background

Most power grids were designed and constructed for a centralized generation system, under the principle of large plants transmitting power over long distances. For grids with increasingly high renewable generators, the intermittent nature of renewable energy production can cause large seasonal swings in the supply [1].

To account for these changes, developing countries often use centralized diesel generation to satisfy energy peak demand each day, as diesel generation is relatively easy to ramp, is not dependent on weather, and has a fairly low fixed cost to add capacity. However, in this study we want to provide a tool to question whether diesel is the most economic way to shave peak demand in some developing countries.

Rwanda is located in East Africa and has significant diesel capacity (28% of national grid generation capacity), but diesel importation is extremely expensive. Most diesel that comes to Rwanda needs to be shipped through either Dar es Salaam or Mombasa, and then trucked overland to Kigali. This results in variable costs of around \$0.38/kwh, depending on current oil prices. Additionally, Rwanda also has rich hydropower resources in the North Province, where Virunga National Park and the Mukungwa River Valley offer highly variable topography and steady rainy seasons.

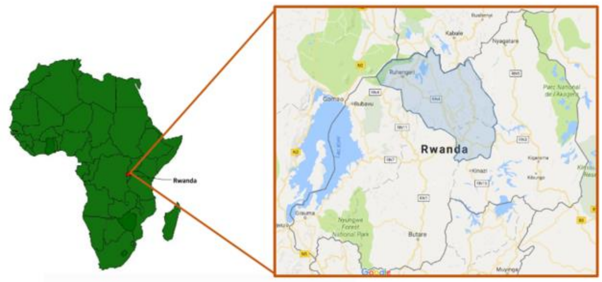

Figure 1 – The North Province of Rwanda holds more than 1000 MW of potential capacity, according to some Rwanda Development Board studies.

In 2016, 43% of Rwanda's installed capacity consisted of hydropower generation [2]. While clean energy constitutes a large portion of Rwanda's current capacity, a central question to this study is whether Rwanda's current hydro resources can be used to economically reduce diesel consumption further.

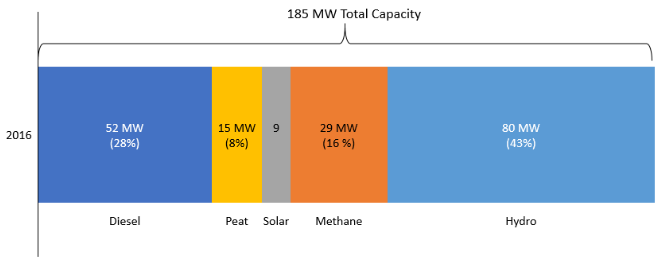

Figure 2 – 2016 Rwanda generation capacity

A majority of Rwanda's installed hydro capacity is run of river hydropower, with little to no storage at the dam. This implies that at any given time, these hydro resources can be capped by available river flow, rather than turbine capacity. It would be possible to add capacity to some of these run of river plants, but at the additional capital cost of lake excavation and weir re-sizing.

For developers, there is no current incentive built into the feed-in tariffs for run of river plants to add storage or peak flow ability to their plants. However, if run of river developers increase storage and turbine capacity, diesel generation would decrease on an annual basis. This occurs because larger storage means that heavy rains can be stored for dry times (rather than sent over the dam, producing no electricity), and higher turbine capacity means that when storage is full and there are heavy flows, the plant can output more power. This study aims to provide a modeling tool with the capability to show the optimal man-made storage and turbine sizing solutions to minimize overall system cost. The essential question being asked is: given current costs of diesel generation and run of river expansion, how much should run of river plants increase their capacity and storage to optimally offset diesel generation? Currently, diesel generation costs the Rwandan Energy Group (REG) around $45 million per year, but this cost would go down if other resources, such as expanded run of river, were implemented.

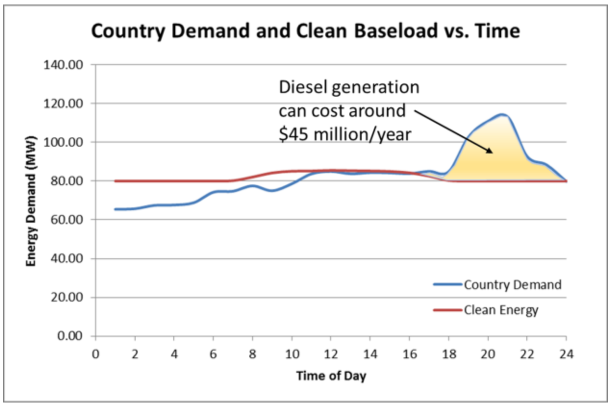

Figure 3 – Diesel is currently being used for peak evening generation, costing around $45 million per year

II. Methodology – Data Collection and Model Overview

For our case study, we used the Mukungwa River Valley, which contains three cascaded hydropower plants. The goal was to characterize the river flow on a daily basis, and to then develop a model of the three run of river plants to determine the energy output of each plant. This energy output would then be fed into a country demand curve, along with the rest of the country's base load, and the difference would be produced with diesel generation. The overall cost of the system would be output. An optimization algorithm would then be run in the background to minimize the cost of the total system, with the constraint that country demand needed to be satisfied.

For the model to have any real accuracy, a seasonal river flow needed to be known, so that the model would have the right river flows to feed into the run of river systems. River flow measurements were taken on the Mukungwa River, and a logarithmic discharge rating curve was used to map discharge measurements to gage measurements. We then mapped this discharge rating curve to eleven years of historical gage measurements, to produce seasonal flows as inputs to our model (Figure 4).

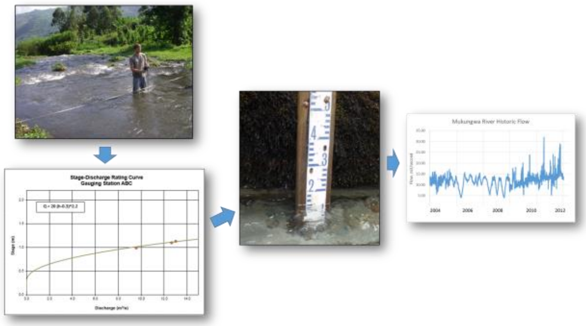

Figure 4 – River discharge measurements were mapped to gage data for seasonal flows.

After river measurements were gathered, a model was developed within MATLAB to determine the overall cost of the system with the constraint of satisfying country demand. The key inputs to this model were:

i. River flow data
ii. Country demand and current baseload
iii. Cost of expanding plant storage and turbine capacity
iv. Precipitation runoff and timing of flows between plants
v. Cost of diesel to cover supply and demand differences

These inputs were fed into the model, and the model then tracks, in a real time simulation:

i. Water in each plant lake
ii. Water flowing between plants
iii. Total energy that needs to be produced by cascaded run of river system

The total energy that needs to be produced by the cascaded system is simply the country demand minus the available baseload. When this does not require the cascaded plants to be producing at 100% of river flow, then water is stored in each plant's lake until the lakes begin to overflow. The decision variables to be adjusted are the storage S_i and the turbine capacity T_i for each plant *i*. The model then calculates the additional cost to the current system in order to meet country demand, and outputs this as a dollar value (Figure 5).

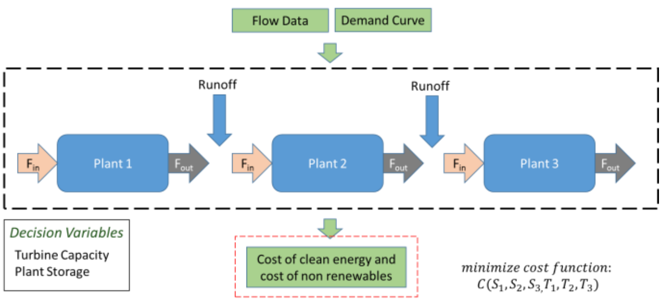

Figure 5 – Pictorial representation of model

Key to this analysis is noting that the objective cost function is actually convex with respect to the storage and turbine values. Each increment of turbine and storage capacity necessarily gives a slightly lower payback than the previous increment, so for a constant cost to increment storage and turbines, this objective cost function is convex.

There are a number of ways to solve convex optimization problems, but for early versions of this model, a simple goal seeking method was used, by adjusting one variable while holding all others constant. This is almost certainly not the most computationally efficient method to use, but for early iterations of the model, the goal was to be sure that the timing of water flows and heuristics of energy production decisions were being accounted for properly [3].

Thus, the basic algorithm for the MATLAB model was as follows:

i. *Begin with a reasonable guess of initial variables S_i and T_i for each plant i.*
ii. *Change one variable according to predetermined step (in our example, it was 1 m^3/s for T_i and 1 hour of storage for S_i.)*
iii. *Run the model for given rainflow to determine overall cost to the system.*
iv. *Check that cost savings are above minimal threshold.*
    a. *If cost savings are below minimal threshold, loop back to (ii), adjusting a different variable.*
    b. *Else, loop back to (ii), adjusting the same variable.*

The general idea is to find the variable that gives the local minimum of the objective function. Once this minimum is found, the algorithm checks to see if adjusting any of the other variables results in minimizing the function further. Once we have a solution where adjusting any of the variables results in an increase of the objective function, we know we've found the global minimum cost.

III. RESULTS

The primary goal of this work was to create a workable model to apply to our case study in the North Province of Rwanda, which was successful. Once the model was complete, the Mukungwa River basin was assessed to look at turbine and storage sizing for the three run of river plants. There are two somewhat interesting capital cost analyses that we wanted to investigate. First, what would have been the optimal storage and plant capacity built were we able to go back to a greenfield state with no storage and no civil works already built? We call this the *Greenfield State.* Second, we ask: what is the best thing to do with the current system? How much more capacity and storage should be added, given the current metrics for all three plants? For instance, Plant 1 already had about 30 days of storage in a lake. This is obviously the minimum storage that can be assessed in the model algorithm, since it doesn't make physical sense to make the lake smaller than it is now. We will address each of these in turn.

A. Greenfield State

We found that, if starting from the greenfield state, an $11.2 million dollar investment into storage and turbine capacity beyond the current state would result in a decrease of diesel production of 62,500 MWH (~$25 million) over a period of five years. $10.5 million of this $11.2 million investment would be for turbine and civil works upgrades, increasing turbine flow capacity by 25% beyond current ability (12 m^3/s). This finding might support the hypothesis that the turbines in place now are sized too small for our given model inputs.

B. Current Mukungwa River Valley System

Additionally, we found that, for the current Mukungwa River Valley system, an additional $7.1 million investment would reduce diesel consumption by 45,000 MWH over 5 years, valued at ~$18 million. Again, much of this investment ($7 million) is for the expansion of flow capacity at the run of river plants, supporting the hypothesis that the turbines are too small for our given model inputs.

C. Feed-in Tariffs for Developers

Our model is also well suited to address the dangers of giving an overwhelming majority of feed-in tariffs to private developers. A short thought experiment follows: one can imagine a scenario where REG offers so many feed-in tariffs that nearly 100% of the national grid has generation incentivized by feed-in tariffs. What will happen to the water in our cascaded system? Developers will produce power whether or not the demand for the power exists, and REG will be held to their fixed contract. REG may request the plants shut down or dump the load, but REG will still be responsible to pay the developers, regardless of the demand at any given time. This will result in water being used in the cascaded system when the energy produced is not necessarily needed.

In a second scenario, imagine that 0% of the national grid is incentivized by private feed-in tariffs, and REG is operating everything cooperatively. In this scenario, REG will only let water flow from the lakes in one of two scenarios: first, if the water will go towards meeting country demand; and second, if the lake is full, and extra spillage is going over the weir. It seems intuitive that much more water would be wasted in the first scenario, by flowing through the system without necessarily meeting energy demand. Our model is well-suited to test this. As an initial check to see if our water volumes were flowing through the system correctly, we compared these two scenarios, of 100% feed-in tariffs to the grid (essentially, developers try to produce as much as possible at all times) and 0% feed in tariffs, where REG only lets water flow when it produces necessary energy. These are displayed in the *High* and *Low* plots shown in Figure 6.

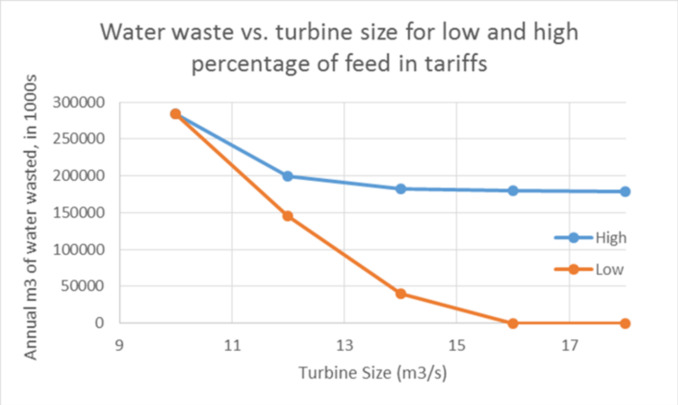

Figure 6 – More water is wasted under high feed-in tariff structures due to non-cooperative operation of plants.

IV. Future Work

While this initial research was successful at developing a model and reliable seasonal rainflow for the Mukungwa River Valley, there is always work of interest that can be furthered. In particular, we see some of the most important work being (1) working with different optimization algorithms, rather than a single goal seeking search, (2) looking at different peak energy costs than standard diesel cost, and (3) allowing for marginal cost curves of storage, rather than a set variable cost.

1. Different optimization Algorithms

Currently, we utilize a relatively simple method of goal seeking for our optimization algorithm. There are much more sophisticated methods, including evolutionary algorithms and genetic algorithms. These could be researched in further detail and assessed for computational efficiency of these sorts of problems.

2. Best opportunity cost

One semi-fatal flaw with the model the way that we have shown it is that it compares run of river capacity increasing with diesel production. This is partially because this is what Rwanda currently does, and also partially because it is low hanging fruit. Almost anything looks preferable to diesel generation from a purely economic viewpoint. However, it is a major assumption that diesel is the best alternative to run of river capacity planning. This may not be the case. For example, good data on pumped storage facilities may show itself to be more economical than the capacity increases that our model suggests. There needs to be a deeper dive into the capital costing of alternative energy technologies other than simple comparing our run of river capacity adjustment to the most expensive form of Rwanda's generation (diesel).

3. Marginal cost curves of storage

This is fairly simple, but our model currently assumes a set variable cost for excavation to make larger lakes at the cascaded plants. However, this may not be the case. One could easily foresee the model allowing for a monotonic curve showing a higher price for each cubic meter of dirt removed than the previous cubic meter. This is not far from being implemented in the model, but is still an improvement from what currently exists.

Acknowledgments

This material is based largely on data that was provided by the Rwanda Energy Group, for which I am very grateful. I am also thankful to the Hydro Research Foundation for sponsoring the work, and Norm Bishop PE, and Dr. Pierre Lermusiaux at MIT for their continued support and always keeping their office doors open and their phones available.

References

[1] KA Folly. "How Much Solar Power Can the Grid Handle?" Africa Utility Week. May 2015, 7-11.

[2] EDCL. "Status of the Hydropower Sector in Rwanda" June 2016

[3] J. Ledolter and R. V. Hogg, Applied Statistics for Engineers and Physical Scientists, Upper Saddle River, New Jersey: Pearson Prentice Hall, 2010