Analysis of Staged Hydraulic Air Compressors

by
Ian Shibley-Styer

A THESIS

submitted to

Oregon State University

Honors College

in partial fulfillment of
the requirements for the
degree of

Honors Baccalaureate of Science in Mechanical Engineering
(Honors Scholar)

Presented March 13, 2023
Commencement June 2023

AN ABSTRACT OF THE THESIS OF

Ian Shibley-Styer for the degree of Honors Baccalaureate of Science in Mechanical Engineering presented on March 13, 2023. Title: Analysis of Staged Hydraulic Air Compressors.

Abstract approved: ____

Deborah Pence

To investigate the viability of staging hydraulic air compressors (HACs), a model for HACs operating with a closed loop water cycle powered by a pump was made. This model was used to simulate how these HACs would respond to variable heights and mass flow ratios for both single-stage and multi-stage applications. Using the results of these simulations and the model, a design tool was made to provide ideal designs for staged HACs when given input, operating, and delivery conditions. The output of this design tool was analyzed to compare single-stage and multi-stage HACs for the same set of delivery conditions.

Keywords: staged hydraulic air compressors, isothermal compression, thermodynamic analysis, optimization.

Corresponding e-mail address: shibleyian@gmail.com

©Copyright by Ian Shibley-Styer
March 13, 2023

Analysis of Staged Hydraulic Air Compressors

by
Ian Shibley-Styer

A THESIS

submitted to

Oregon State University

Honors College

in partial fulfillment of
the requirements for the
degree of

Honors Baccalaureate of Science in Mechanical Engineering
(Honors Scholar)

Presented March 13, 2023
Commencement June 2023

Honors Baccalaureate of Science in Mechanical Engineering project of Ian Shibley-Styer presented on March 13, 2023.

APPROVED:

____

Deborah Pence, Mentor, representing Mechanical, Industrial and Manufacturing Engineering

____

James Liburdy, Committee Member, representing Mechanical, Industrial and Manufacturing Engineering

____

Ivan Nepomnyashchikh, Committee Member, representing Mechanical, Industrial and Manufacturing Engineering

____

Toni Doolen, Dean, Oregon State University Honors College

I understand that my project will become part of the permanent collection of Oregon State University, Honors College. My signature below authorizes release of my project to any reader upon request.

____

Ian Shibley-Styer, Author

Table of Contents

Nomenclature

g	gravity, m s^{-2}
H	height, m
$\dot{m}$	mass flow rate, kg s^{-1}
P	pressure, Pa
P^*	pressure, dimensionless
R	gas constant, J kg^{-1} K^{-1}
T	temperature, K
v	velocity, m s^{-1}
$\dot{W}$	power, W
y	fraction of air not absorbed
z	elevation, m

Greek Letters

η	efficiency
ρ	density, kg m^{-3}

Subscripts

a	air
D	downcomer
R	riser
w	water
m	air and water mixture
ratio	water/air
P	pump

1. Introduction and Background

Hydraulic air compressors (HACs) are mechanisms that compress air using the downward flow of water [1]. Figure one shows a schematic for an example HAC. Water is collected tens of meters above the application height, at what are considered the inlet conditions, and then allowed to flow with gravity in a contained manner [2]. As the flow develops, the Venturi effect induces a lower pressure inside the downcomer shaft, which pulls air from the surroundings into the water, creating a bubbly mixture [3]. As these bubbles travel down, increasing hydrostatic pressures compress the air before it reaches the separation chamber and is extracted for use [4]. Historically, HACs have been used in mining applications since the technology was developed by Charles Taylor in 1890. 18 HACs have been reported, the largest was in Ontario Canada and supplied up to 25 silver mines with compressed air during its 70 years of near continuous operation [2].

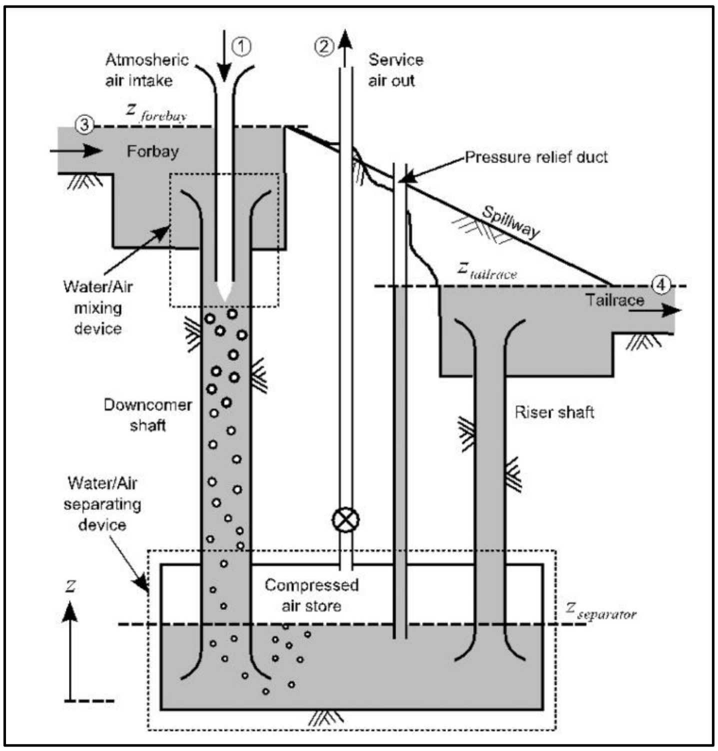

Figure 1. A generalized design for a hydraulic air compressor [1]

Hydraulic air compressors have several advantages over the traditional mechanical versions. HACs operate nearly isothermally as opposed to traditional gas compressors that can only be said to operate nearly adiabatically [4]. A truly isothermal process is the theoretical ideal for compression work. HACs operate close to this ideal because the larger mass flow rate and specific heat value of water give it a higher thermal capacitance than the air [3]. Additionally, the air out of the separation chamber is dryer and cooler than traditional compressors, and it is often cleaner because it is free from machine oil or other contaminants [5]. Dryer air is better for many applications because the water vapor in moist air can freeze during depressurization. Cooler air can also be more useful depending on the application. As a result of these benefits, HACs could soon become more common in industrial practices [2].

One area where hydraulic air compressors have a growing appeal is the energy sector. They are a method of compressing air that does not require the use of fossil fuels, and compressed air has the potential to be a valuable large scale energy storage technique [5]. Compressed air can be combined with gas turbines to produce power in a Brayton cycle [4]. Traditionally, this cycle is done using staged mechanical compressors, which heat the air significantly during each compression phase, resulting in efficiency loses. Compressed air is also a compliment to renewable energy sources because it can be used to store energy and balance load demands [6]. Coupling air compressors with technology like wind turbines allows excess energy from windy periods to be saved for moments of high electricity demand. The lack of these storage capabilities has been an obstacle to the replacement of fossil fuels in energy production. Researchers are also looking at HAC applications in wave and tidal energy production [5]. These would not harness water flows in the same way historic HACs have, but the constant supply of water and the energy potential of waves could prove useful. Overall, the possible applications of HACs are growing rapidly as people reconsider the technology, and many of these uses could be important to renewable energy.

Due to the potential for hydraulic air compressors to become common use installations, understanding how to optimize their efficiency is an important area of research. It was determined that the ratio of the mass flow rates of water and air would be critical in understanding the relationship between the compression work out and the hydraulic work in [4]. Further research and analysis of historical HACs determined that there was an optimal mechanical efficiency based on mass flow ratios and delivery pressures, but it was noted that the reported efficiencies of many past HAC systems were likely overestimates because they did not account for air absorption during the compression [1]. The exact mass flow ratio is difficult to control because it depends on the entrainment mechanism [4]. However, it is important that the mass flow of water is at least three orders of magnitude greater than that of air so the flow can be considered isothermal [1]. Additional research has determined gas entrainment varies based on the velocity of the water and the geometry of the eductor. This work observed entrainments rates that equate to a flow ratio range of 1000 to 1700 [7].

Peet-Mati Sööt from CMM Energy LLC has proposed that hydraulic air compressors could be staged to allow for their use in industrial applications that do not already have a large elevation change. These systems would pump water up to the top of the downcomer shaft in a closed loop design. The use of a pump could also be beneficial in expanding HAC potential beyond requiring a natural water flow. For example, these systems could be used on offshore wind turbines. As has been discussed prior, excess power from turbines could be used to compress air, and a pumped HACs could be an efficient way of accomplishing that goal. Additionally, the limitless supply of ocean water would be available for use in a pumped system. These ideas remain theoretical for now, but the introduction of pumps and stages greatly expands the applications potential for HACs. The efficiency of staged HACs will be different than historic installations, so it is important to discover what will impact their performance. To examine this, a model has been designed to test how simulated HACs would respond to changes in design parameters.

2. Model Design

To determine the efficiency of a staged hydraulic air compressor, the compression work done will be compared to the pump work required to drive the system. Figure 2 shows the simplified HAC design used for this analysis. It is assumed that the temperature and pressure of the air is equal to that of the water at the inlet.

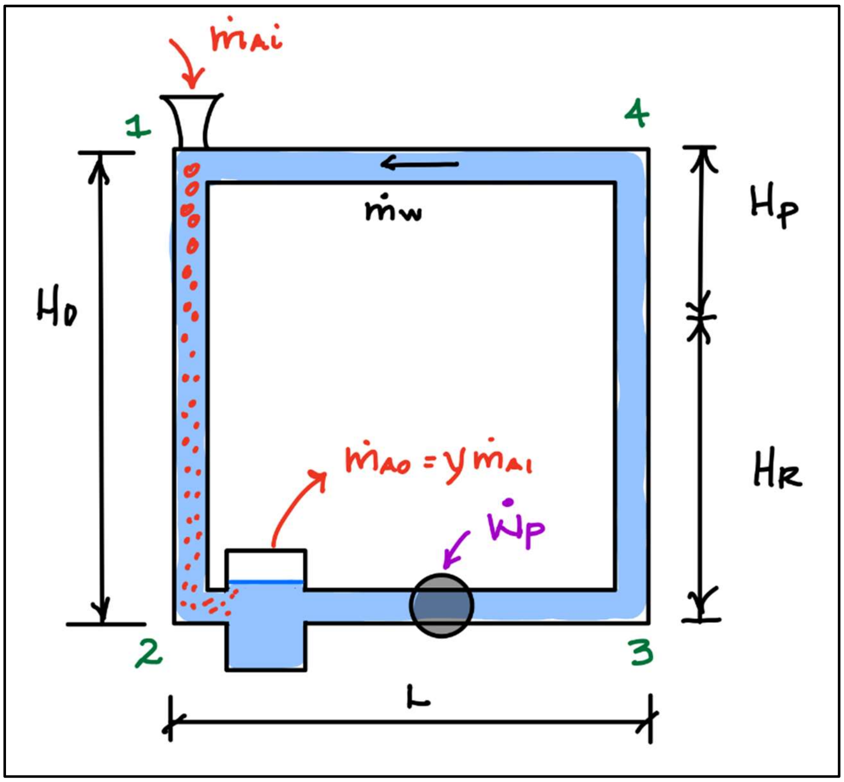

Figure 2. A simplified closed loop HAC

The pipe is assumed to be large enough that friction losses can be ignored. Based on this assumption, the energy equation can be written in a simplified form.

$$\frac{P_1}{g\rho_1} + \frac{v_1^2}{2} + z_1 = \frac{P_2}{g\rho_2} + \frac{v_2^2}{2} + z_2 \qquad (1)$$

Neglecting any change in velocity along the downcomer, the energy equation can be combined with the ideal gas equation of state to determine the temperature of the air at the bottom of the downcomer, T_{a2}. T_{a1} is the air temperature at the inlet, H_D is the height of the downcomer, and R_a is the gas constant for air.

$$T_{a2} = T_{a1} + \frac{gH_D}{R_a} \tag{2}$$

The water pressure at the bottom of the downcomer, P_{w2}, is found using the energy equation, assuming no change in velocity. P_{w1} is the water pressure at the inlet and ρ_m is the average mixture density.

$$P_{w2} = P_{w1} + \rho_m g H_D \tag{3}$$

The average mixture density across the downcomer, ρ_m, is assumed to be the average of the mixture density at the inlet, ρ_{m1}, and the mixture density at the bottom of the downcomer, ρ_{m2}.

$$\rho_m = \frac{\rho_{m1} + \rho_{m2}}{2} \tag{4}$$

The mixture density for any point along the downcomer, ρ_{mj}, can be found by a manipulation of the conservation of mass. ρ_{aj} is the density of air at a point along the downcomer and can be found using the ideal gas law and assuming the pressure of the air is always equal to the pressure of the water. ρ_w is the density of water.

$$\rho_{mj} = \frac{1 + \dot{m}_{ratio}}{\frac{1}{\rho_{aj}} + \dot{m}_{ratio}\frac{1}{\rho_w}} \tag{5}$$

The mass flow rate ratio, $\dot{m}_{ratio}$, is defined as the mass flow rate of water, $\dot{m}_w$, over the mass flow rate of air, $\dot{m}_a$.

$$\dot{m}_{ratio} = \frac{\dot{m}_w}{\dot{m}_a} \tag{6}$$

The difference in density between the downcomer shaft and the riser means hydrostatic effects do not cause the water to return to the elevation the downcomer inlet, H_D, but instead to an intermediate height, H_R. The pressure at H_R is equal to the pressure at H_D, so the pump head required to close the loop, H_p, can be expressed only in terms of H_D and densities.

$$H_P = (1 - \frac{\rho_m}{\rho_w}) H_D \tag{7}$$

The pump power required to operate a single stage, $\dot{W}_P$ can be found using H_P. η_P is the isentropic pump efficiency.

$$\dot{W}_P = \frac{\dot{m}_w g H_P}{\eta_P} \tag{8}$$

The coefficient of performance, η, can be found by taking the isothermal compression work over the pump work. y is the percent of air not absorbed into the water during compression.

$$\eta = \frac{y \dot{m}_a R_a T_{a1} \ln\left(\frac{P_{a2}}{P_{a1}}\right)}{\dot{W}_P} \tag{9}$$

It was also of interest to look at a measure of the change in pressure, which was non-dimensionalized for comparisons.

$$P^* = \frac{y \dot{m}_a (P_{a2} - P_{a1})}{\rho_{a1} \dot{W}_P} \tag{10}$$

A computer simulation was created to understand how H_D and $\dot{m}_{ratio}$ affect these measures of performance in a single stage and across multiple stages. This code can be found in Appendix A.

3. Results

To better understand how the downcomer height, H_D, and the mass flow ratio, $\dot{m}_{ratio}$, affect η and P^*, both equations can be rewritten.

$$\eta = \frac{yR_aT_{a1}\ln\left(\frac{P_{a2}}{P_{a1}}\right)}{\dot{m}_{ratio}H_D\frac{g\left(1-\frac{\rho_m}{\rho_w}\right)}{\eta_P}} \tag{11}$$

$$P^* = \frac{y(P_{a2}-P_{a1})}{\dot{m}_{ratio}H_D\frac{g\left(1-\frac{\rho_m}{\rho_w}\right)}{\eta_P}\rho_{a1}} \tag{12}$$

For all simulations, the inlet temperature and pressure were set as 300 K and 100 kPa respectively, the density of water was assumed to be 1000 kg/m^3, the acceleration of gravity was assumed to be 9.81 m/s^2, the isentropic pump efficiency is assumed to be 0.85, and the percent of air not absorbed into the water was assumed to be 0.95.

The performance metrics were investigated for a single stage before moving to multiple stages. Figure 3 shows a plot of P^* vs downcomer height at four different mass flow rate ratios. For each $\dot{m}_{ratio}$ tested, P^* always increases with increasing H_D. This is because larger values of H_D cause a greater hydrostatic pressure at the bottom of the downcomer, which results in a greater pressure difference. Looking at the equation for P^*, there is an H_D term in the denominator from diving by pump work, but the change in pressure in the numerator must be more impactful as the plot indicates a positive relationship between height and P^*. The $\dot{m}_{ratio}$ plots cross when H_D is about 35 meters, meaning that height is where having a lower $\dot{m}_{ratio}$ goes from resulting in a lower P^* to resulting in a higher P^*. A lower $\dot{m}_{ratio}$ results in a lower mixture density, which primarily affects the overall pressure change in the downcomer and the pump head in the riser. A lower flow ratio results in a lower pressure change and a greater pump head, both of which would expect to result in a lower P^*. However, the plot suggests there is a more influential interaction between H_D and $\dot{m}_{ratio}$ that results in lower a P^* for larger values of H_D at larger values of $\dot{m}_{ratio}$.

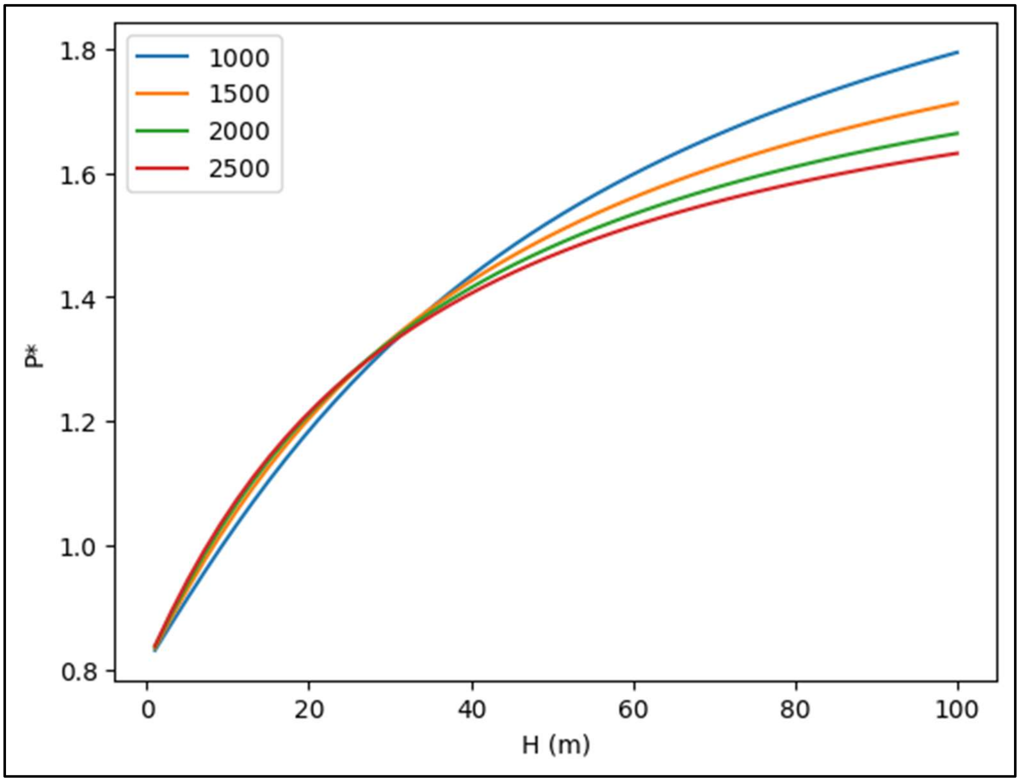

Figure 3. P^* vs H_D at varied $\dot{m}_{ratio}$

Figure 4 shows a plot of η vs H_D at four different values of $\dot{m}_{ratio}$. For all the simulated mass flow ratios, η decreases as height increases. This is because increasing H_D will increase the pump work needed to close the loop at a faster rate than increasing the isothermal compression work done by the system. This can be seen by examining the equation for η. The pump head needed to close the loop will be a specific fraction of the downcomer height that is determined by $\dot{m}_{ratio}$, so the denominator will increase in a mostly linear fashion. It can also be seen that the compression work in the numerator will increase in a logarithmic fashion. The result of the ratio of these terms is the negative relationship seen in the plot. All of the plots converge on 0.81 as height goes to zero. This is because the assumed value of y and the isentropic pump efficiency remain constant while the other variables remove themselves in this edge case. The product of y and the isentropic pump efficiency is 0.8075.

There were additional interesting observations about this plot. There is an inflection point in each line, which means there is a downcomer height where the size of the increase in η from decreasing H_D begins to diminish. This could be useful information for optimization. Additionally, higher mass flow rates seem to be impacted more by changes in height. This is likely a result of the higher mixture density having a greater overall effect when it acts across larger heights. It is important to note that $\dot{m}_{ratio}$ and H_D do not work independent of each other. The trend of performance characteristics being dependent on interactions between H_D and $\dot{m}_{ratio}$ was seen in Figure 3 as well. Finally, Figure 4 shows the relationship between η and H_D works opposite the relationship between P^* and H_D. The most efficient systems come from the smallest

H_D values and therefore have the smallest overall pressure changes, while the larger heights result in the greatest pressure changes and the least efficient systems.

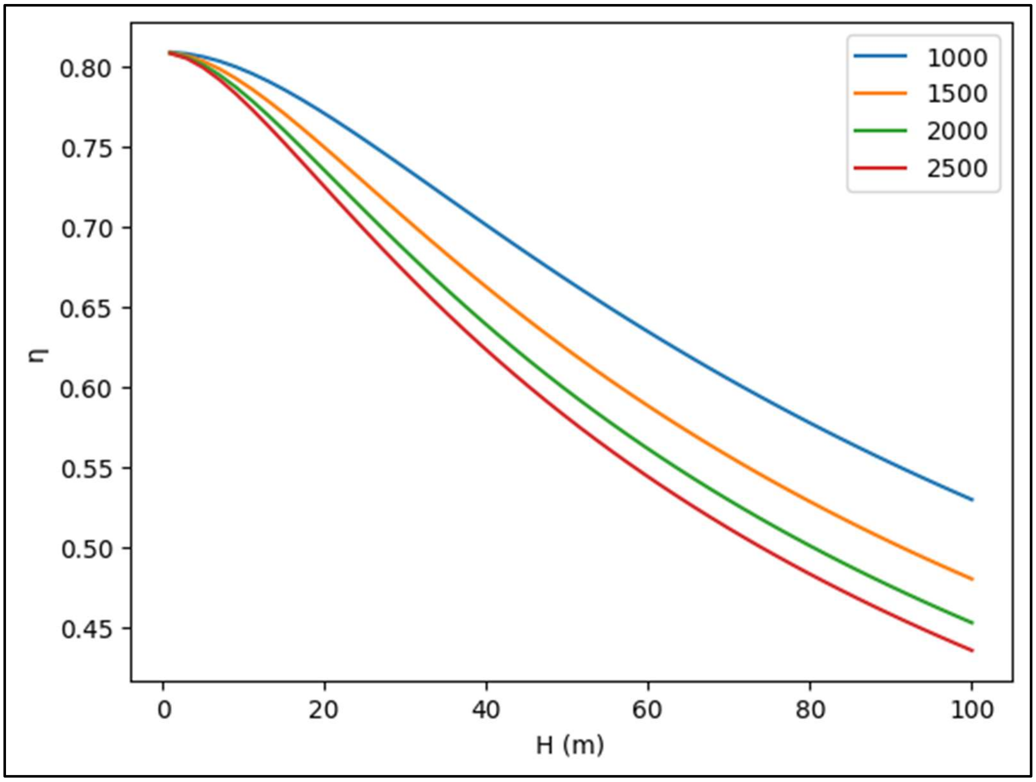

Figure 4. η vs H_D at varied $\dot{m}_{ratio}$

After observing the opposing trends in the relationship between H_D and the two performance metrics, it was of interest to consider a combination η and P^* that could indicate an optimal balance between pressure increase and system efficiency. Figure 5 shows a plot of ηP^* vs H_D at four values of $\dot{m}_{ratio}$. The plot shows a peak H_D for each $\dot{m}_{ratio}$ that could be considered optimal. Figure 4 shows that the lower flow ratios have inflection points at higher values of H_D, which supports the idea that HACs with low $\dot{m}_{ratio}$ values are optimized at larger heights. This could be because the lower mixture density needs more height to develop the optimal pressure change. Looking at the equations, the peak is likely a result of the interaction between P^* increasing from the mostly linear pressure difference term and η decreasing as the logarithmically increasing pressure ratio term is outpaced by the pump work. The lower mixture density would appear to cause this interaction to develop slower and peak later.

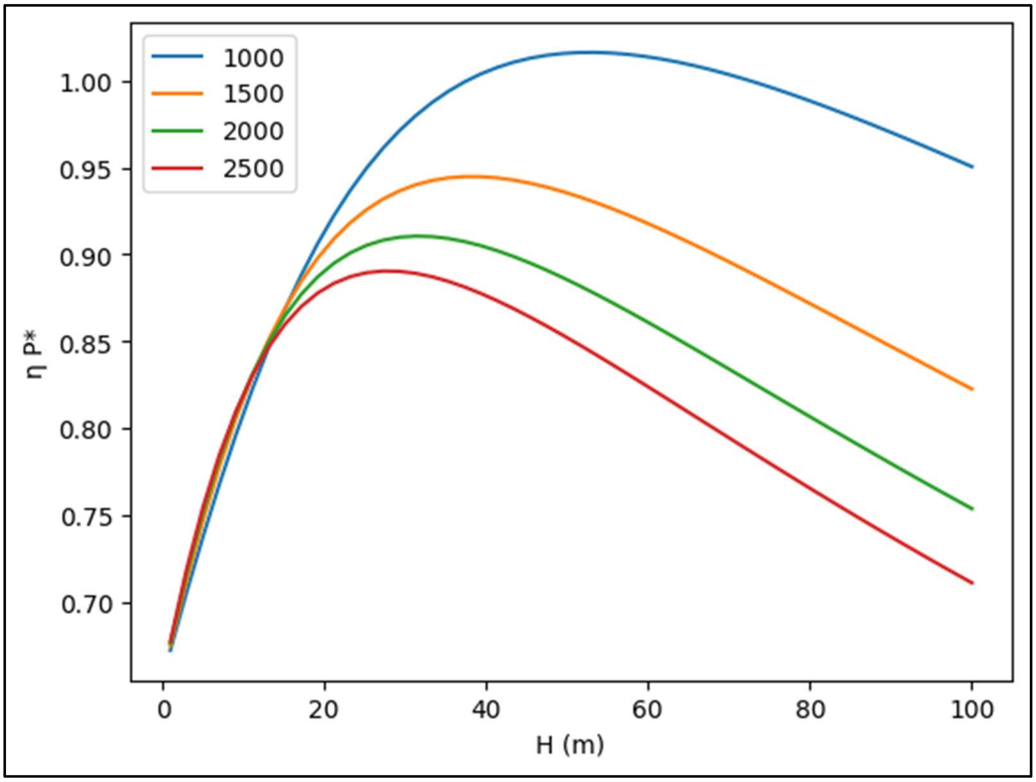

Figure 5. ηP^* vs H_D at varied $\dot{m}_{ratio}$

With an understanding of how to select optimal heights, the next area of focus was how adjusting the value of $\dot{m}_{ratio}$ would affect η. Figure 6 shows a plot of η vs $\dot{m}_{ratio}$ for three different H_D values in meters. The plot shows that η decreases as $\dot{m}_{ratio}$ increases. This is likely a result of less air being available to be compressed at higher flow ratios. As the ratio increases, ρ_m goes to ρ_w, and the mixture becomes closer to pure water. Less air being compressed gives a lower overall compression work. The $\dot{m}_{ratio}$ term in the denominator represents this relationship because it will drive η to zero as it increases. This also means that η will increase with lower flow ratios. However, there will be a physical limit to this trend because at some much smaller value of $\dot{m}_{ratio}$, buoyancy forces would overcome the drag from the water and air would escape the stream. This phenomenon is not captured well by the model but is not of great concern because HACs are limited by the amount of air they can entrain via the eductor. Effectively, the mass flow ratio is already within the physical limits of the air water interaction for real systems. To account for this, the analysis was restricted to ranges where ratios were seen as reasonable based on prior findings.

Also of interest, the mass flow ratio has more of an effect when it is over a larger height. This is likely a result of the pressure term because that is what depends on both H_D and $\dot{m}_{ratio}$, a relationship that was observed during the analysis of H_D. Again, this is caused by the fact that $\dot{m}_{ratio}$ affects the density which can then have a greater impact on delivered pressure with a greater height.

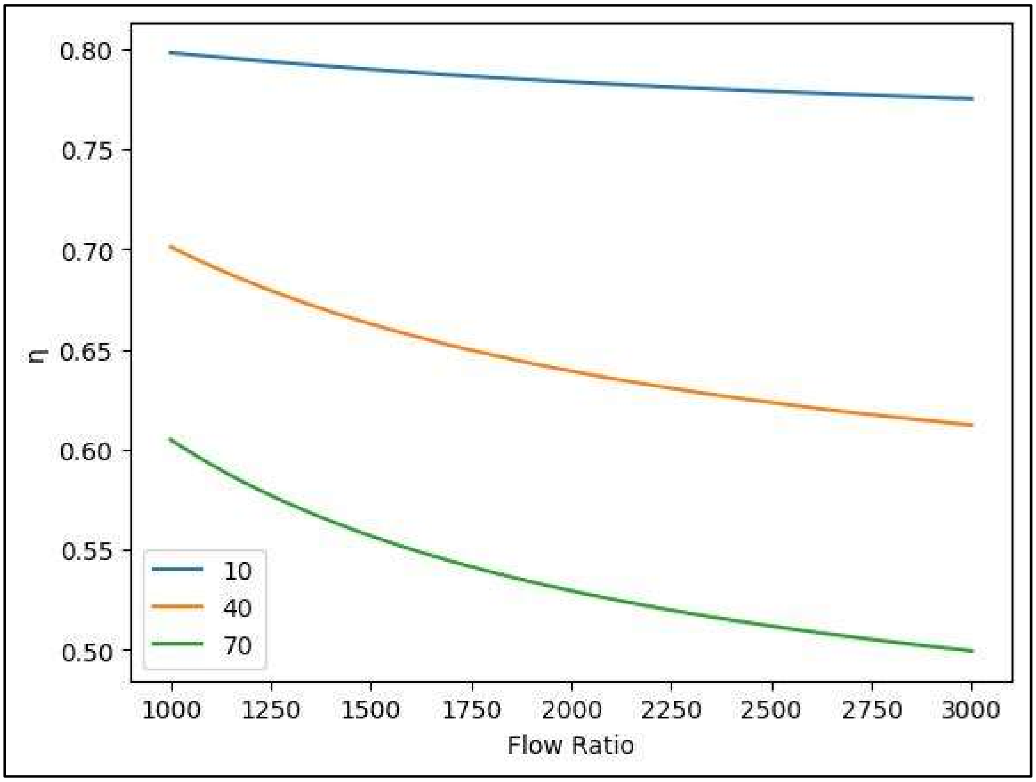

Figure 6. η vs $\dot{m}_{ratio}$ at varied H_D (m)

After understanding how η was affected by varied flow ratios, the next question was how η was affected by staging. Figure 7 shows η vs $\dot{m}_{ratio}$ over multiple stages. For this plot, H_D was held at 10 meters for each stage. Each subsequent stage inlet is set to the outlet conditions of the $\dot{m}_{ratio}$ that resulted in the highest η for the prior stage. The major notable pattern is η increasing for each subsequent stage. The primary difference between stages is starting pressure, meaning the process is more efficient when operating at higher pressures. This could be because a higher inlet pressure means a higher mixture density with close to the same amount of air flow. When looking at the equation for η, driving the density ratio in the denominator to one would drive the entire denominator to zero. An opposite trend would result from the pressure ratio in the numerator also going to one at higher inlet pressures, which drives the log of this value to zero, but this must happen at a slower rate. The plot indicates the higher density is more beneficial for η overall.

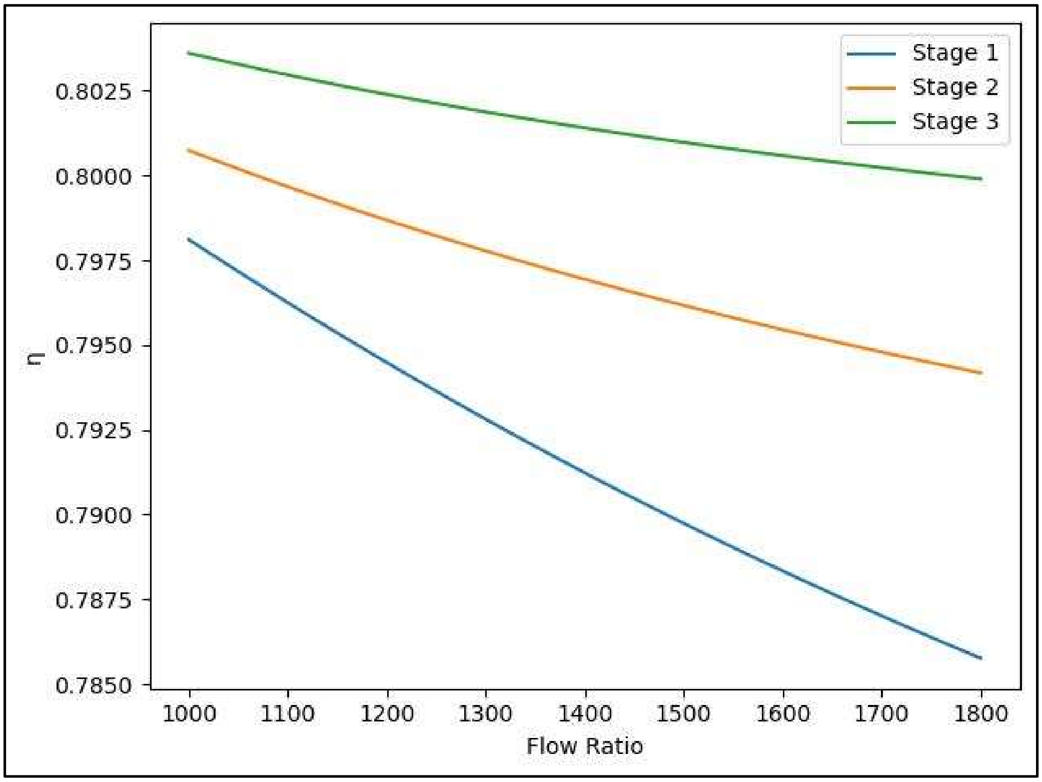

Figure 7. η vs $\dot{m}_{ratio}$ over consecutive stages

After discovering the positive effect staging has on η, there was interest in creating a design tool that takes a desired mass flow rate of air and delivery pressure and returns the ideal configuration for a staged hydraulic air compressor system. This tool also takes inlet conditions, a maximum height for the stages, and a range of possible mass flow ratios.

Figure 8 shows the relationship between the number of stages required for an HAC system to reach set delivery conditions and the average η and total pump power for each separate system. By examining this plot, it is clear that the average η improves with additional stages while the total power required decreases. This would indicate it is better to operate pumped HACs in a staged manner because it is more efficient and uses less power overall. It is important to note the marginal benefit of each additional stage decreases quickly. This means the economics and logistics of additional stages will likely determine how many stages would be necessary for a specific application, but for most, a handful of stages will be ideal. This graph was made using a series of results from the design tool that were sorted based on the number of stages they required. The highest η and corresponding pump work were selected for each set of results with the same number stages required. This means the delivery conditions were held constant, but the reported mass flow ratios and heights are variable from one point on the plot to another. This plot was meant to be an overall summary of results from comparing design tool outputs and a more detailed explanation of the process that created it can be found below.

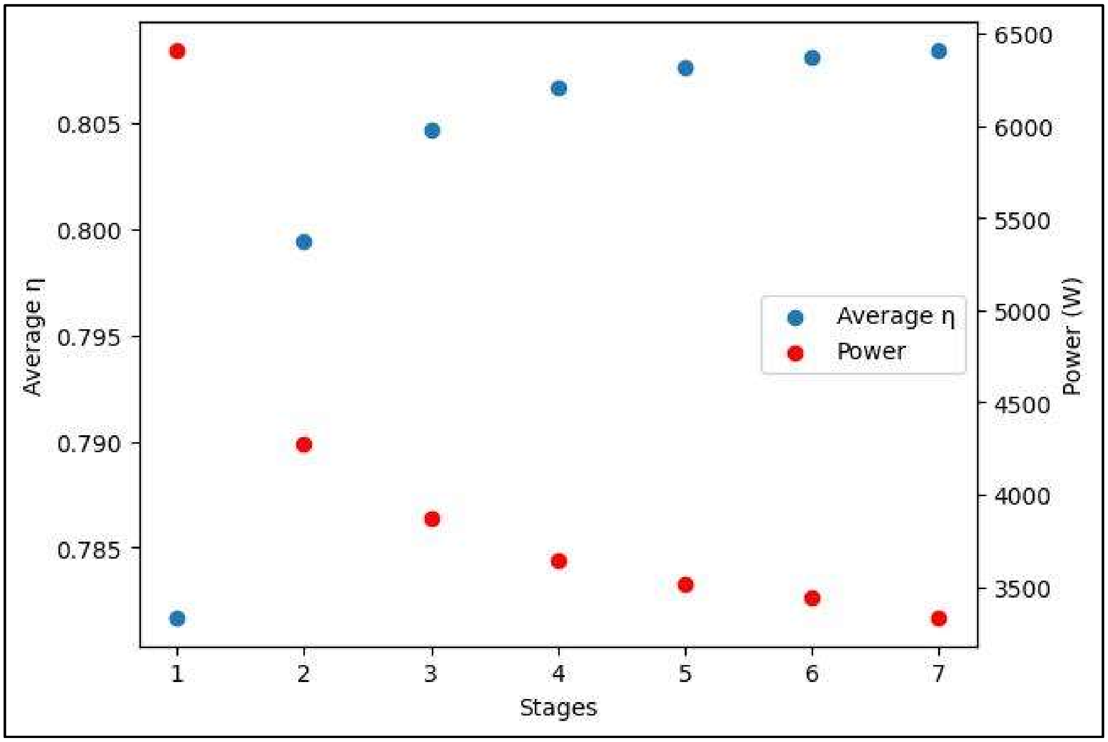

Figure 8. The best average η and corresponding total pump power for each set of design tool outputs that had an equal number of stages required.

The design tool works by creating a range of height values from 2 meters to the given max and then breaking this and the mass flow ratio range into separate sets of evenly spaced values. It then runs these vectors through the HAC simulation that was developed. This outputs a series of arrays for values like pressure out and η. The tool then looks through the array and throws out any values that do not meet the delivery pressure requirement. If no values meet the delivery conditions, it will select the highest pressure output, add an additional stage, and repeat this process. It was decided to select the highest pressure output when delivery condition are not met because analysis of staged HACs indicated higher pressures were more efficient. Once it finds one or more configurations that can meet the delivery conditions, it will compare the values of η and select the highest one. The tool then returns the number of stages required, the mass flow rate of air required into the first stage, the optimal configuration of flow rate and height for each stage, and the pressure, η, and pump work values resulting from each stage.

In order to explore how staging effects η, the design tool was used to generate results for a series of inputs that only varied the max allowable height. The desired output pressure and flow rate were 200 kPa and 0.2 kg/s respectively. The range of mass flow ratios given was 1000 to 2000. The inlet conditions were 300 K and 100 kPa.

Figure 9, 10, and 11 show a graphical summary of the design tool results for a max height of 15, 10, and 5 meters respectively. Figure 9 shows a single stage that ran at an $\dot{m}_{ratio}$ of 1325, had an η of 0.776 and required 4880 W to power the pump. Figure 10 shows a system that uses two stages. The first is 10 meters tall, operates at an $\dot{m}_{ratio}$

of 2000, had an η of 0.783 and required 2545 W. The second is 4 meters tall, operates at an $\dot{m}_{ratio}$ of 1000, had an η of 0.808 and required 1268 W. Figure 11 shows a system that required three stages. The first is 5 meters tall, operates at an $\dot{m}_{ratio}$ of 2000, had an η of 0.801 and required 1379 W. The second is 5 meters tall, operates at an $\dot{m}_{ratio}$ of 2000, had an η of 0.804 and required 1117 W. the third is 3.86 meters tall, operates at an $\dot{m}_{ratio}$ of 1020, had an η of 0.808 and required 1224 W.

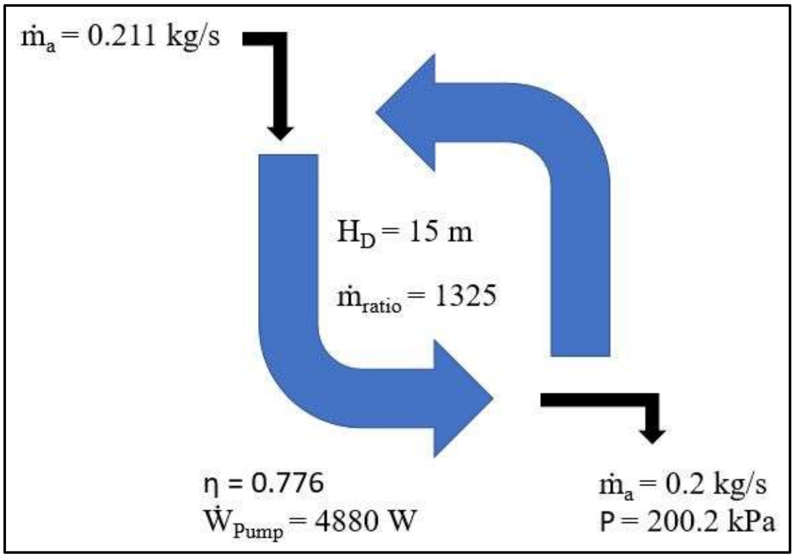

Figure 9. The design tool results for a max height of 15 m

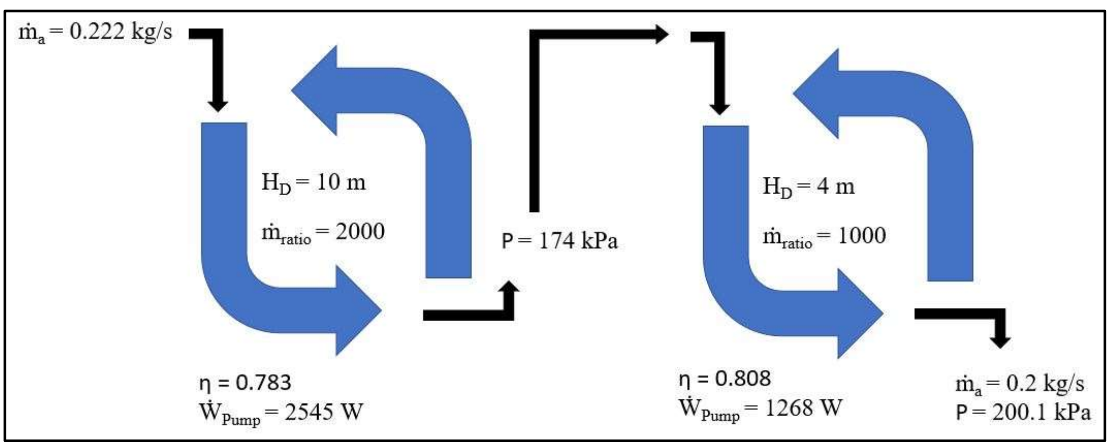

Figure 10. The design tool results for a max height of 10 m

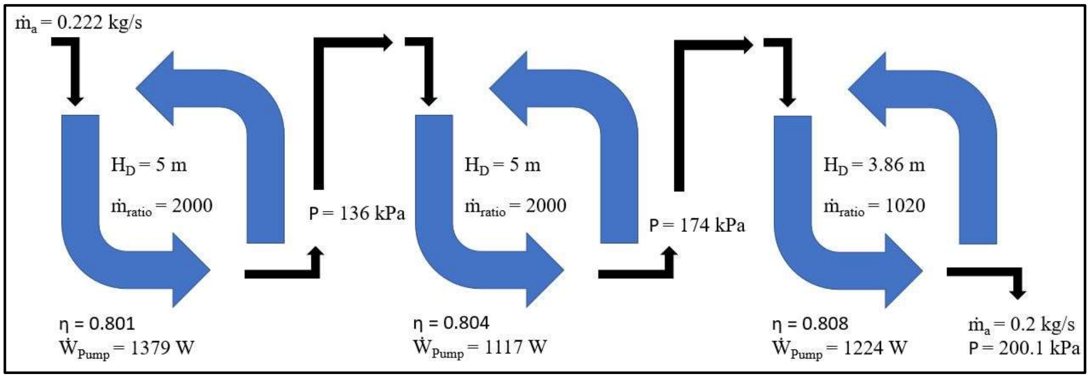

Figure 11. The design tool results for a max height of 5 m

By examining the design tool returns for these specific instances, it is clear that the η for each stage increases when there are multiple stages in the system. A trend towards lower total pump power can also be seen. At this point, it was beneficial to understand how average η and total pump power respond to variation in max allowable height, and how these values correlated to the number of stages required.

Figure 12 shows the stages required and average η for a series of design tool requests with different max heights. Each point represents an HAC system determined by the tool, just like those represented more in depth in Figures 9-11. The line shows how many stages are needed in those systems. When examining the point at 10 meters, it appears that 2 stages are needed and the average η is about 0.797, which agrees with Figure 10. Overall, the efficiencies seem to be improving with additional stages. The reason there is an increase in average η before a second stage is added is because the tool selects the best value at each stage, not the best overall average. Another important observation is the drop in η just before a second and third stage is added.

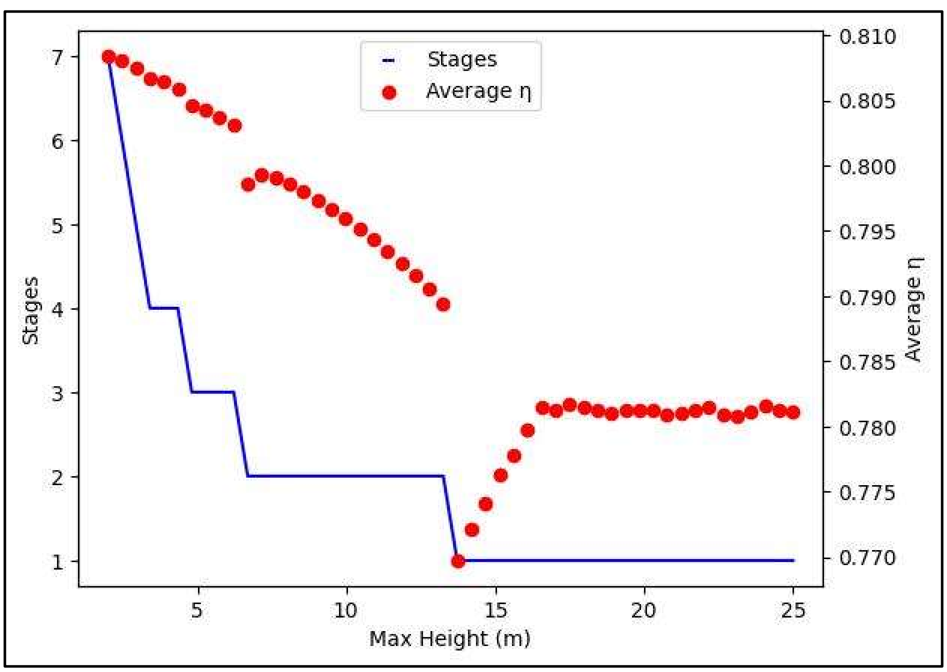

Figure 12. The design tool results for stages required and average η vs the max allowable height

Figure 13 shows the stages required and total pump power for a series of design tool requests with different max heights. The total power varies inconsistently within each category of required stages. This is mostly because pump power itself is not being prioritized, so the tool's selection process does not create a smooth pattern for this particular output. However, there is a drop in total power used at the height just before an additional stage is needed, similar to the pattern seen in Figure 12.

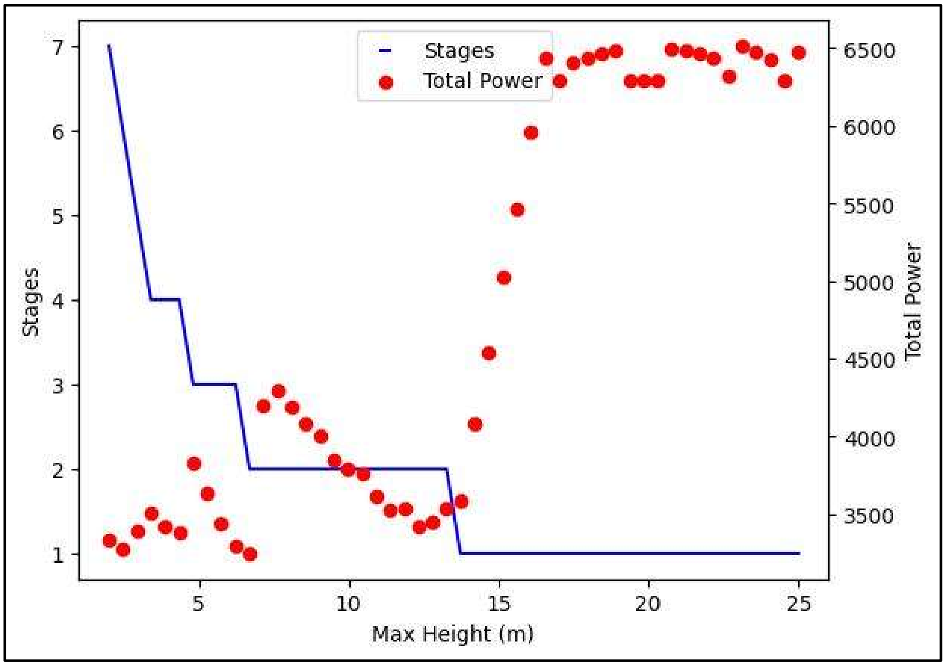

Figure 12. The design tool results for stages required and total power vs the max allowable height

Investigation of the phenomenon occurring just before an additional stage was added showed that it is a result of how the tool selects HAC configurations. As the maximum allowable height gets closer to the cutoff that will require an additional stage, there reaches a point when the desired pressure can be reached only with an inefficient configuration because better ones are now too tall. This forces the tool to pick a worse configuration because it has no other options.

The variability of the average η without adding a stage makes it difficult to determine the overall pattern of η with respect to stages. However, all these results can be grouped based on the number of stages required, and a best η can be selected from each grouping. This was the process used to generate Figure 8.

4. Conclusion

The results of the single-stage analysis of hydraulic air compressors showed that η decreases and P^* increases with increasing H_D. The product, ηP^*, was used to determine the ideal height for an HAC operating at a given flow rate. It was also found that η decreases with increasing $\dot{m}_{ratio}$. The results of multi-stage analysis showed that η increases with subsequent stages.

These results were used to guide the development of a design tool that could take input and delivery conditions and return an optimized HAC system. Analysis of the output of this tool for a variety of maximum allowable heights was used to determine that overall, staging HACs will increase the average η and decrease the total pump power required for a system. This suggests staged HACs could be viable alternatives to the traditional design.

Future work should be done to explore how variable absorption rates might affect staged systems, especially if the water is recycled continuously. Further work could also be done to explore how the geometry of the piping and the design of the eductor might influence the ideal configurations. Additionally, work could be done to find other ways to represent the efficiency over multiple stages beyond an average and optimize the output of a single stage beyond ηP^*. Finally, the design tool was built primarily for analysis, improvements could make it more user friendly for actual design.

Appendix A: Code

This appendix includes the Python code used to simulate individual and staged hydraulic air compressors and the design tool that was created to select configurations. It also includes the code used to make the graphs found in the results section.

```
import numpy as np
import matplotlib.pyplot as plt

# Constants
g = 9.81  # gravity in m/s^2
density_water = 1000  # In kg/m^3
R = 287  # J/(kg*K)


def hac_output(test_heights, mass_flow_ratio, diameter=0.15, velocity=5,
temp_i=300, pressure_i=100000,
            ise_pump_eff=0.85, y=0.95):
    """
    hac_output takes a vector of test downcomer heights and mass flow ratios ([flow of
water]/[flow of air]) and returns
    the final temperature, final pressure, pressure change from inlet to outlet,  a non-
dimensional pressure change, a
    coefficient of performance, pump work, mass flow of water, and mass flow of air
for a single stage. It can also take
    pipe diameter, water velocity, inlet temperature and pressure, isentropic pump
efficiencies, and values of y (the
    percent of air no absorbed into water by Henry's law).

    :param test_heights: a numpy array of downcomer heights to be tested
    :param mass_flow_ratio: a numpy array of mass flow ratios to be tested
    :param diameter: optional pipe diameter to determine mass flow of water (m).
    :param velocity: optional water velocity to determine mass flow of water (m/s)
    :param temp_i: inlet temperature (K)
    :param pressure_i: inlet pressure (Pa)
    :param ise_pump_eff: isentropic pump efficiency
    :param y: percent of air not absorbed into water by henry's law
    :return: final temp and pressure, pressure difference, non-dimensional pressure
difference, coefficient of
    performance, pump work, mass flow rate of water, mass flow rate of air.
    """
    # pre define variables for later code
    pressure_f = 0
    average_density = 0
    pressure_check = np.ones(len(test_heights))
```

```
    # reshape test_heights array into 1xn row vector
    test_heights = np.reshape(test_heights, (1, len(test_heights)))
    # reshape mass_flow_ratio into a nx1 column vector
    mass_flow_ratio = np.reshape(mass_flow_ratio, (len(mass_flow_ratio), 1))

    # find the cross-sectional are of the pipe (m^2)
    area = np.pi / 4 * diameter ** 2
    # calculate mass flow rates of water and air (kg/s)
    mass_flow_water = density_water * area * velocity
    mass_flow_air = mass_flow_water/mass_flow_ratio

    # inlet pressure is initial guess for final pressure (Pa)
    pressure_f_start = pressure_i
    # find density of air at inlet using ideal gas laws (m^3/kg)
    density_air_1 = pressure_i / (R * temp_i)
    # find temperature final of air using ideal gas law (m^3/kg)
    temp_f = temp_i + g * test_heights / R
    # calculate the density of the air and water mixture (using conservation of mass?)
(m^3/kg)
    density_mix_1 = (1 + mass_flow_ratio) / (1 / density_air_1 + mass_flow_ratio /
density_water)

    # iterative loop to find the final pressure
    while pressure_check.any() > 0.01:
        # find the density of air at the bottom of the downcomer using assumed final
pressure (m^3/kg)
        density_air_2 = pressure_f_start / (R * temp_f)
        # find the density of the mixture at bottom of downcomer (m^3/kg)
        density_mix_2 = (1 + mass_flow_ratio) / (1 / density_air_2 + mass_flow_ratio /
density_water)
        # find the average of the mixture density using the densities at the inlet and
bottom of the downcomer (m^3/kg)
        average_density = (density_mix_2 + density_mix_1) / 2
        # find final hydrostatic pressure using the average density (Pa)
        pressure_f = pressure_i + test_heights * g * average_density
        # check if new calculated pressure matches old guess and update final pressure
to new value
        pressure_check = abs(pressure_f - pressure_f_start)
        pressure_f_start = pressure_f
    # find height difference between inlet and riser shaft water height from hydrostatic
effects (m)
    pump_head = (1 - average_density / density_water) * test_heights
    # work done by pump to overcome pump (W)
    pump_work = mass_flow_water * g * pump_head / ise_pump_eff
    # compute coefficient of pressure as isothermal compression work over pump work
```

```
    hac_efficiency = y * R * temp_i * np.log(pressure_f / pressure_i)/(mass_flow_ratio
* g * pump_head / ise_pump_eff)
    # find pressure change of air from inlet to outlet (Pa)
    pressure_diff = pressure_f-pressure_i
    # multiply pressure difference by volumetric flow rate of air at the inlet and divide
by pump work
    non_dim_pressure_diff =
mass_flow_air*pressure_diff*y/(density_air_1*pump_work)

    return [temp_f, pressure_f, pressure_diff, non_dim_pressure_diff, hac_efficiency,
pump_work, mass_flow_water,
            mass_flow_air]


def staged_hac(max_height, mass_flow_range, pressure_out, flow_air_out, t_i, p_i):
    """
    staged_hac is a function to determine how many stages are needed to meet a given
pressure and air flow rate out
    when given a set of restrictions about height and mass flow ratios.
    :param max_height: The max height of a given stage
    :param mass_flow_range: the range of mass flow ratios to test across
    :param pressure_out: the desired final pressure of the air
    :param flow_air_out: the desired mass flow rate of air
    :param t_i: inlet temperature
    :param p_i: inlet pressure
    :return: returns the number of stages, the required air intake rate at the first stage,
and lists of mass flow
    ratios, heights, pressures, etas, and pump work values for each stage
    """
    # set aside variables for use in output
    h_values = []
    mr_values = []
    pressure_values = []
    eta_values = []
    pump_work_values = []

    # create arrays of height and mass flow ratios to test
    h = np.linspace(1, max_height, 50)
    m_r = np.linspace(mass_flow_range[0], mass_flow_range[1], 50)

    # start at stage 1
    stage = 1

    # Loop while the pressure into the next stage (p_i) is less than the requested
pressure out
    while p_i < pressure_out:
```

```
        # run a simulation for the given conditions
        sim = hac_output(h, m_r, temp_i=t_i, pressure_i=p_i)

        # if any output pressures are greater than the requested pressure out, set aside all
required information
        if np.any(sim[1] > pressure_out):
            # create an array of the eta values
            etas = sim[4]
            # set any eta value equal to zero if the pressure is not above the requested
pressure out
            etas[sim[1] < pressure_out] = 0
            # find the index of the highest eta value from the reduced list
            best_eta_index = np.unravel_index(etas.argmax(), etas.shape)
            # append output lists with appropriate values
            mr_values.append(m_r[best_eta_index[0]])
            h_values.append(h[best_eta_index[1]])
            eta_values.append(etas[best_eta_index])
            pressure_values.append(sim[1][best_eta_index])
            pump_work_values.append(sim[5][best_eta_index])
            # set p_i to a pressure higher than the out pressure so the loop will terminate
            p_i = sim[1][best_eta_index]
        # if no pressure is high enough
        else:
            # find the highest pressure output from the simulation array results
            best_p_index = np.unravel_index(sim[1].argmax(), sim[1].shape)
            # set conditions to the output of the highest pressure out simulation
            t_i = sim[0][0][best_p_index[1]]
            p_i = sim[1][best_p_index]
            # append appropriate values for output
            mr_values.append(m_r[best_p_index[0]])
            h_values.append(h[best_p_index[1]])
            pressure_values.append(sim[1][best_p_index])
            eta_values.append(sim[4][best_p_index])
            pump_work_values.append(sim[5][best_p_index])
            # add a stage
            stage = stage + 1
            # while loop will repeat because outlet pressure is not high enough
    # calculate the air in required assuming a y value of 95 for each stage
    air_in = flow_air_out/(0.95**stage)
    # return parameters
    return [stage, air_in, mr_values, h_values, pressure_values, eta_values,
pump_work_values]


def stage_analysis():
    # graph eta vs mass flow ratio at h = 4 over three stages
```

```
    h = np.array([10])
    m_r = np.linspace(1000, 1800, 100)
    x = m_r
    stages = 3
    p_i = 100000
    t_i = 300
    for i in range(stages):
        sim = hac_output(h, m_r, temp_i=t_i, pressure_i=p_i)
        y_cop = np.reshape(sim[4], len(x))
        plt.plot(x, y_cop, label=('Stage ' + str(i+1)))
        index = np.argmax(y_cop)
        t_i = sim[0]
        p_i = sim[1][index]
    plt.legend()
    plt.xlabel('Flow Ratio')
    plt.ylabel('\u03B7')
    plt.tight_layout()
    plt.savefig('Stage.jpg', bbox_inches='tight')


def efficiencies_vs_height():
    # graph height vs eta, P_star and eta*P_star
    h = np.linspace(1, 100, 50)
    m_r = np.linspace(1000, 2500, 4)
    x = h
    string_list = ['1000', '1500', '2000', '2500']
    for i in range(4):
        string = '$' + string_list[i] + '$'
        sim = hac_output(h, np.array([m_r[i]]))
        y_cop = np.reshape(sim[4], len(x))
        y_non_dim_p = np.reshape(sim[3], len(x))
        plt.plot(x, y_cop*y_non_dim_p, label=string)
    plt.xlabel('H (m)')
    plt.ylabel('\u03B7P*')
    plt.legend()
    plt.savefig('COP_P_Star_vs_H', bbox_inches='tight')


def efficiencies_vs_flow_ratio():
    # graph flow ratio vs eta and P_star
    h = np.linspace(10, 70, 3)
    m_r = np.linspace(1000, 3000, 100)
    x = m_r
    string_list = ['10', '40', '70']
    for i in range(3):
        string = '$' + string_list[i] + '$'
```

```
        sim = hac_output(np.array([h[i]]), m_r)
        y_cop = np.reshape(sim[4], len(x))
        y_non_dim_p = np.reshape(sim[3], len(x))
        plt.plot(x, y_cop*y_non_dim_p, label=string)
    plt.xlabel('Flow Ratio')
    plt.ylabel('\u03B7 P*')
    plt.legend()
    plt.savefig('COP_P_Star_vs_Flow_Ratio.jpg', bbox_inches='tight')


def stages_vs_power():
    """
    This function runs a series of simulations for a set of conditions with a variable
max allowable height. It takes
    the result of these simulations and groups them categorically based on the number
of stages required to meet outlet
    conditions. It them graphs the number of stages vs the maximum possible eta from
all configurations that result in
    a given number of stages and the corresponding total power required.
    :return:
    """
    # define array of max heights to test
    max_h = np.linspace(2, 50, 500)
    # define arrays for use later
    stages_needed = np.zeros(len(max_h))
    total_power = np.zeros(len(max_h))
    average_eta = np.zeros(len(max_h))

    # get simulation results for each max height
    for i in range(len(max_h)):
        # for each max height value run the simulation to determine the most optimal
configuration
        sims = staged_hac(max_h[i], [1000, 2000], 200000, 0.2, 300, 100000)
        # save the stages needed, total power and average of the etas across stages for
each simulation
        stages_needed[i] = sims[0]
        total_power[i] = sum(sims[6])
        average_eta[i] = np.average(sims[5])

    # convert stages to integer type variable to be usable in inequality later
    stages_needed = stages_needed.astype(int)
    # create empty arrays for graph values
    max_etas = np.zeros(max(stages_needed))
    min_powers = np.zeros(max(stages_needed))
    x = np.zeros(max(stages_needed))
```

```
# loop through each set of heights requiring the same number of stages
for i in range(max(stages_needed)):
    # create copies of arrays to save original data
    temp_power = total_power.copy()
    temp_eta = average_eta.copy()
    """if the simulation did not return the number of stages being considered in this
iteration of the loop set the
    power value to infinity and the eta value to zero. This negates all values not
associated with a set number
    of stages"""
    temp_power[stages_needed != i+1] = np.inf
    temp_eta[stages_needed != i + 1] = 0
    # find the max eta and the index where it occurs for each set of heights with a
given stage number requirement
    max_etas[i] = max(temp_eta)
    index = np.argmax(temp_eta)
    # save the corresponding total power required
    min_powers[i] = temp_power[index]
    # make the x-axis for plotting
    x[i] = i+1
# graph the average eta and total power vs stages needed
fig, ax1 = plt.subplots()
ax2 = ax1.twinx()
ln1 = ax1.scatter(x, max_etas, label='Average \u03B7')
ax1.set_xlabel('Stages')
ax1.set_ylabel('Average \u03B7')
ln2 = ax2.scatter(x, min_powers, c='r', label='Power')
ax2.set_ylabel('Power (W)')
lns = [ln1, ln2]
labs = [ln.get_label() for ln in lns]
ax1.legend(lns, labs, loc=7)
plt.savefig('Total_Power_vs_Stages.jpg', bbox_inches='tight')
```

Citations

[1] Pavese, V., Millar, D., and Verda, V., 2016, "Mechanical Efficiency of Hydraulic Air Compressors," Journal of Energy Resources Technology, **138**(6): 062005.

[2] Millar, D., 2014, "A review of the case for modern-day adoption of hydraulic air compressors," Applied Thermal Engineering, **69**(1), pp. 55-77.

[3] Young, S., 2017, "Simulating air absorption in a hydraulic air compressor (HAC)," M.S. thesis, Natural Resources Engineering, Laurentian University.

[4] Bidini, G., Grimaldi, C.N., Postrioti, L., 1997, "Thermodynamic analysis of hydraulic air compressor-gas turbine power plants," Proceedings of the Institution of Mechanical Engineers, **211**(5), pp. 429-437.

[5] Roy, A. K. and Jumar, K., 2020, "Performance Evaluation of Low Head Hydraulic Air Compressor: A Prospective Source of Renewable and Green Energy," Advanced Materials Manufacturing & Characterization, **10**(1), pp. 14-28.

[6] Kim Y.M., Shin, D.G., Favrat, D., 2011, "Operating characteristics of constant-pressure compressed air energy storage (CAES) system combined with pumped hydro storage based on and energy and exergy analysis," Energy, **36**(10), pp. 6220-6233.

[7] Yadav, R.L., Patwardhan, A.W., 2008, "Design aspects of ejectors: Effects of suction chamber geometry," Chemical Engineering Science, **63**(15), pp. 3886-3897.