EPSRC
Engineering and Physical Sciences Research Council

InFoMM
Industrially Focused Mathematical Modelling

EPSRC Centre for Doctoral Training in Industrially Focused Mathematical Modelling

A novel approach to cost-effective energy extraction in rivers and tides

Graham Patrick Benham

UNIVERSITY OF OXFORD

VerdErg
Renewable Energy

Contents

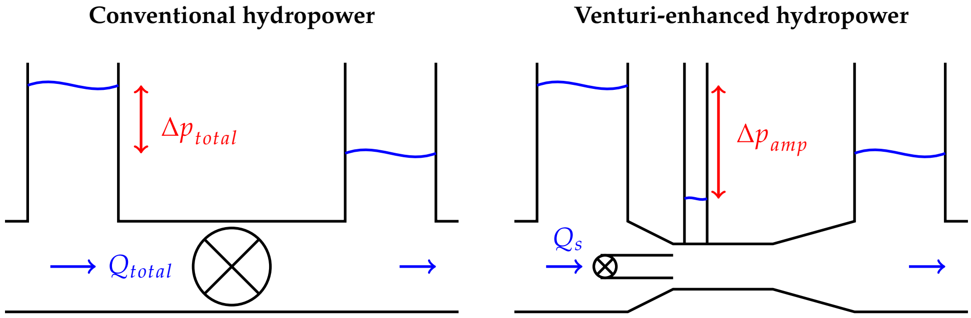

Figure 1 – Comparison between conventional and Venturi-enhanced hydropower. In conventional hydropower a pressure drop Δp_{total} drives a flow rate Q_{total} across a turbine. In Venturi-enhanced hydropower a reduced flow rate Q_s passes through the turbine, whilst the remaining flow is accelerated through a contraction, amplifying the pressure drop across the turbine to Δp_{amp}. The illustrated levels of water are indicative of the pressure (high level corresponds to high pressure and vice versa).

1 Introduction

1 The economics of low hydropower

Conventional high head hydropower is the most cost effective means for renewable energy generation. However, the cost-effectiveness of conventional hydropower becomes progressively impaired in low head situations of 3 m or less. Because power output is a function of head and flow, low head hydropower projects require a large volumetric flow rate to be economically viable. The resulting technology is a large, slow moving turbine, which requires a gearbox and huge civil installations to house it. The increased cost of such equipment erodes the economics of those installations.

Head refers to the change in water height across a wier, or a tidal barrier. Dams can be anything between hundreds of meters tall (high head) and just a few meters (low head).

We study a novel type of hydropower, pioneered by VerdErg Renewable Energy, which produces cost effective power from low head sources by amplifying the head drop. The majority of the available flow is squeezed through a narrow contraction (known as a *Venturi*), thereby increasing the pressure that a turbine in the remaining minority of the flow experiences. It therefore allows that a low pressure/high flow project location is treated as a high pressure/low flow project. With this approach, the turbine used is small and compact, requiring less civil infrastructure to be housed, and will rotate at higher speed, avoiding the need for a gearbox. These changes result in substantially reduced project costs, and therefore cost of energy. It allows projects to be realised at locations that have previously been considered uneconomical or have become uneconomical due to recent political targets. In addition, this approach is attractive for its low environmental impact due to the fact that aquatic life can pass through the turbine-free primary flow unharmed.

The economic improvements compared to conventional technologies make this technology attractive to investors, landowners and communities with access to low-head sites across the world. In England and Wales alone 25,935 potential sites have been identified, which, if developed, could provide 1,178 MW of power. In addition, globally there is a low head hydropower potential of around 98 GW. In terms of tidal sites, in the UK there is a potential of 45GW, which is nearly entirely untapped, and so is the worldwide potential of 571 GW. In the long term, this technology will aid the reduction of carbon emissions, the reduction of the cost of energy, and the improvement of the security of energy.

1 The hydrodynamics of low head hydropower

In Figure 1 we compare conventional hydropower to *Venturi-enhanced* hydropower. In conventional hydropower (a) a pressure drop[1] Δp_{total} drives a flow rate Q_{total} across a turbine. The total available power is

$$P_{total} = \Delta p_{total} \times Q_{total}. \tag{1}$$

[1]Note that Δp_{total} here refers to the overall drop in static pressure, rather than the drop in 'total pressure', which is sometimes used in the literature to refer to the sum of the static pressure and the dynamic pressure (kinetic energy per unit volume).

In low head[2] situations, where Δp_{total} is small (a few metres, perhaps), the conventional approach may not be cost-effective, as described earlier.

Instead, we scale down the size and capacity of the turbine using a Venturi contraction. In Venturi-enhanced hydropower (b) a reduced portion of the total flow (secondary) with flow rate Q_s passes through a pipe with constant cross-section containing the turbine, whilst the majority of the flow (primary) is diverted around the turbine and accelerated through a contraction. According to the Bernoulli principle, the static pressure drops as the flow accelerates, such that the Venturi contraction amplifies the pressure drop across the turbine to Δp_{amp}. The use of the Venturi effect here is similar to a mechanical gearing system. Instead of a turbine dealing with the full flow and a low head, it deals with a reduced flow and an amplified head, thereby allowing for cheaper electricity production.

The Bernoulli principle implies that if the velocity of a fluid u increases, the pressure p drops, and vice versa. Mathematically it states that $p + \frac{1}{2}u^2$ =constant.

Similarly to (1), the total power generated by Venturi-enhanced hydropower is

$$P_V = \Delta p_{amp} \times Q_s. \tag{2}$$

It is important to note that (1) and (2) are not necessarily equal. In particular, by separating the flow into two parts, accelerating one part, and mixing them back together again, there is a fundamental energy loss (resulting in a loss of power). This is due to viscous dissipation associated with shear stress between the flows as they mix together, which we will discuss later. The ratio between (2) and (1), which we call the hydrodynamic efficiency, is given by

$$\eta = \frac{\Delta p_{amp} Q_s}{\Delta p_{total} Q_{total}}. \tag{3}$$

Equation (3) can be rewritten in terms of two key parameters, the flow rate fraction $\mathcal{F} = Q_s/Q_{total}$, and the pressure ratio $R_p = \Delta p_{amp}/\Delta p_{total}$, such that

$$\eta = \mathcal{F} \times R_p. \tag{4}$$

In order to improve the cost-effectiveness of Venturi-enhanced hydropower, we balance keeping $\mathcal{F}$ small, reducing the cost of the turbine, whilst keeping R_p as large as possible, increasing the generated power. However, R_p itself, and hence the efficiency η, depend on $\mathcal{F}$ as well as other factors (such as the shape of the Venturi), and this dependence is unknown. Therefore, we focus on understanding how the efficiency depends on the hydrodynamics in the Venturi, and in turn how to design the Venturi to maximise efficiency.

1 Glossary of terms

- **Flow rate:** The volume of water passing through a pipe per second.
- **Pressure drop:** The negative change in pressure measured between two points in a pipe.
- **Pressure recovery:** The positive change in pressure measured between two points in a pipe.
- **Head drop:** Equivalent to pressure drop.
- **Low/high head:** Small/large pressure drop.
- **Venturi contraction:** A narrow contraction in a conduit which accelerates the flow and lowers the pressure.
- **Bernoulli principle:** A law in fluid dynamics relating the pressure and velocity.
- **Primary flow:** The part of the flow which does not pass through the turbine, accelerated through a contraction.
- **Secondary flow:** The part of the flow which passes through the turbine unaccelerated.
- **Shear layer:** A thin high-stress region of fluid between two streams of different speeds

[2]The head drop is equal to the pressure drop divided by ρg, where ρ is the density and g is the gravitational acceleration.

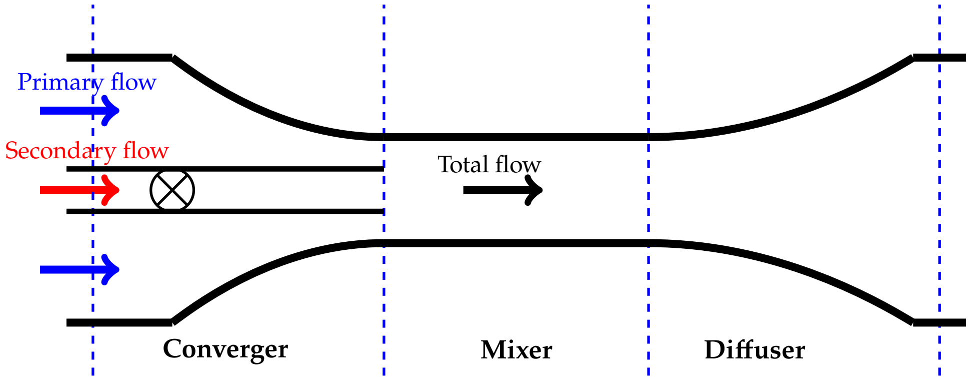

Figure 2 – Schematic diagram showing a cylindrical cross-section of Venturi-enhanced hydropower. The primary flow is accelerated through a Venturi contraction (converger), amplifying the pressure drop across the secondary flow, which flows through the turbine. Both flows combine in the mixer to form the total flow, before expanding and decelerating in the diffuser.

- **Turbulence:** A state of flow, characterised by unsteadiness, chaotic motion and eddies.
- **Boundary layer:** A thin high-stress region of fluid near a solid wall
- **Boundary layer separation:** A scenario where the boundary layer near the solid wall grows dramatically, no longer thin, and may contain regions of flow reversal.
- **Uniform/non-uniform flow:** A uniform flow has constant velocity across its cross-section. A non-uniform is flow anything else, such as the confluence of the primary and secondary flows.
- **Flow development:** The manner in which a flow transforms from being non-uniform to uniform as it travels downstream.

1 The converger, mixer and diffuser

In Figure 2 we illustrate an example of Venturi-enhanced hydropower. The geometry is cylindrical and is divided into a converger, a mixer and a diffuser. The function of each of the three sections is as follows:

- **Converger**: A contracting annular section of pipe which accelerates the primary flow. The secondary flow, including the turbine, is kept in a straight pipe of constant cross-section running along the same axis as the primary flow. By the Bernoulli principle, the pressure decreases as the primary flow is accelerated, such that the turbine feels an amplified pressure drop.
- **Mixer**: A straight section of pipe in which the primary and secondary flows mix together. Although there is significant mixing, the flow is not expected to be uniform or fully developed by the end of the mixer.
- **Diffuser**: A widening section in which the flow decelerates and the pressure increases.

The Kutta condition and Boundary Layer Theory are two theoretical tools developed in the early 20th century to understand aircraft design. We applied them here to Venturi-enhanced hydropower.

The primary function of the converger is to amplify the pressure drop across the turbine. In order for this to work, the low pressure induced by the acceleration of the primary flow must be felt by the secondary flow where they meet, at the end of the converger section. In other words, the pressure must be uniform across the end of the converger section. One of our first results was proving this to be true using a combination of *Boundary Layer Theory* and the *Kutta condition*. We validated our theoretical explanation with experimental measurements and computational fluid dynamics simulations.

The inflow of the mixer is composed of the slower secondary flow in the centre and the faster primary flow on the outside. The primary function of the mixer is to enable the two flows to exchange momentum (and energy), producing a combined flow which is as uniform as possible. Of the whole geometry, the mixer has the narrowest cross-section, producing the strongest wall drag. Therefore, the length of the mixer must be chosen

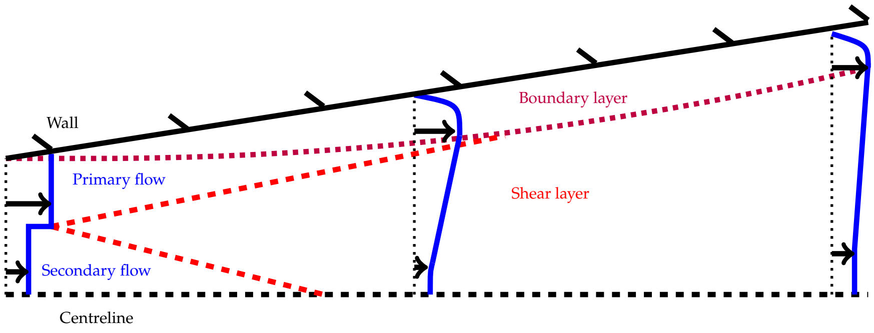

Figure 3 – Schematic diagram in half a cylindrical cross-section of the mixer/diffuser, illustrating the layer structure of the flow, which is the basis for our mathematical model. Velocity profiles are illustrated at three locations. The primary and secondary flow regions have distinct uniform velocities, and are separated by a turbulent shear layer where they mix together, and which has a linearly varying velocity. There is also a boundary layer near the channel wall.

carefully, such that it is long enough that the primary and secondary flows mix well, but not so long that wall drag causes undue energy losses. Besides energy losses due to wall drag, there is also an inevitable energy loss due to viscous dissipation associated with shear stress between the flows as they mix together.

The diffuser section has the function of converting high speed low pressure flow to low speed high pressure flow. Diffuser performance is characterised by *pressure recovery*, which is a measure of the pressure gained between the inflow and the outflow, relative to the kinetic energy at the inflow. Since the pressure at the diffuser outflow is fixed (determined by the hydrostatic pressure of the river/weir), the pressure recovery of the diffuser affects the inflow pressure only. Therefore, a diffuser with large pressure recovery has a relatively low inflow pressure. Since pressure is continuous between the mixer and diffuser, a low pressure at the diffuser inflow contributes to the pressure amplification across the turbine. Therefore, increasing diffuser pressure recovery causes a direct increase in the power generated by the turbine.

A velocity profile (blue lines in the adjacent plot) is a method of plotting the velocity across a region of flow, with velocity on the x axis and distance on the y axis.

Since the inflow of the diffuser is not fully mixed, the diffuser is partly responsible for mixing the flow, similarly to the mixer. Therefore, although the current design distinguishes the mixer and the diffuser as separate sections, we find it useful to consider the mixer and the diffuser as the same section, where the combined function is to mix the primary and secondary flows together, whilst simultaneously recovering as much pressure as possible. Since convergers are well-studied in the literature, we focus our research on the mixer/diffuser.

2 Mixing flows inside a channel

Within the mixer/diffuser, the primary and secondary flows mix together and expand. There are several fluid dynamical aspects that play an important role in this process, and must be taken into account when creating a mathematical model:

The Reynolds number of a flow Re measures the relative importance of inertial and viscous forces. We consider a high Re (turbulent) regime, where viscous forces are almost undetectable, except at very small scales.

- **Fundamental energy loss**: Due to the fact that the flows have different speeds, there is a fundamental energy loss associated with their mixing, corresponding to viscous dissipation (*sticking* the flows together).
- **Fundamental pressure rise**: Accompanying the fundamental energy loss, as the flows mix together, there is also a corresponding fundamental rise in pressure. This is a consequence of conservation of momentum.
- **Wall drag**: In addition to the fundamental energy loss, there will also be some additional energy lost due to friction at the channel walls.
- **Turbulent flow**: Since the Reynolds number for this flow is very high ($Re \approx 10^6$), the flow is in the turbulent regime, characterised by unsteady, chaotic motion. This makes it very difficult to predict exact instantaneous flow behaviour, so instead it is typical to describe the *mean* flow behaviour.

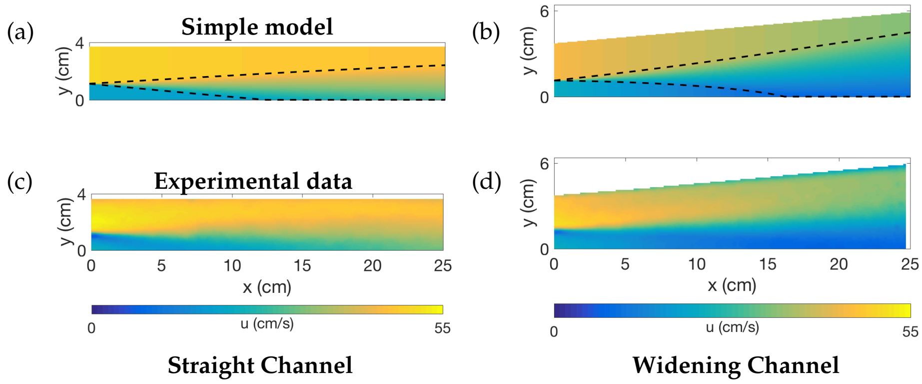

Figure 4 – Comparison between our model and (scaled down) experimental data for a straight channel (a, c), and a widening channel (b, d). Flow is from left to right and is symmetric about the channel centreline so we only show half of the channel. Colour plots show the time-averaged velocity in the x direction, u. In the case of the simple model, we also overlay dashed lines representing the width of the shear layer.

- **Shear layers**: When the primary and secondary flows meet, since they both have very different speeds, they mix together in a violent turbulent pattern of flow, called a shear layer. The way in which these shear layers form and develop is strongly dependent on the shape of the channel, amongst other factors, and have a large impact on efficiency.
- **Boundary layers**: There are also regions of flow near the channel walls which undergo large amounts of stress, called boundary layers. If the channel expands too rapidly, these may grow dramatically with drastic consequences on efficiency. This is known as *boundary layer separation*.

In Figure 3 we sketch an illustration of the flow in the mixer/diffuser, showing the various different phenomena at play. By dividing the flow into distinct regions, as illustrated, we created a mathematical model to describe the time-averaged spatial development of the flow, serving as an intuitive and low-computational-cost predictive tool. We validated this model by comparison with *Particle Image Velocimetry (PIV)* experiments. As shown in the velocity colour plots in Figure 4, there was very close agreement. Hence, the next step was to use it for design and optimisation purposes.

Particle image velocimetry (PIV) is an experimental technique that uses particle tracking to measure the velocity field of a region of flow.

3 Shape optimisation of the channel

Having determined the important phenomena that occur inside the mixer/diffuser, and having created a predictive mathematical model to describe the flow, we formulate an optimisation to determine the channel shape that will provide maximum power. As is conventional, the optimisation problem is broken down into an objective (what is being optimised), decision variables (what can be modified) and constraints (what are the rules).

3 Objective

Upon careful analysis, it transpires that the maximum power generated is achieved when the channel is designed to give maximal *pressure recovery*. Pressure recovery is a measure of the pressure gained in the mixer/diffuser, relative to the kinetic energy flux at the inlet. Since the downstream and upstream pressures are determined by the river or tide, the pressure recovery in the mixer/diffuser determines the operating point of the turbine, and hence the power.

Therefore, we set our design objective as the mass-averaged pressure recovery coefficient, which is defined as

$$C_p = \frac{\int_{\text{outlet}} \text{Pressure flux} - \int_{\text{inlet}} \text{Pressure flux}}{\int_{\text{inlet}} \text{Kinetic energy flux}}, \tag{5}$$

where the pressure and kinetic energy fluxes are given by integrating across the inlet and outlet cross-sectional areas.

	Pros	Cons
Shallow angles	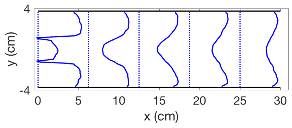 Good flow development (Efficient momentum exchange)	Large energy loss due to enhanced drag from narrow walls
Wide angles	Small energy loss due to reduced drag from wide walls	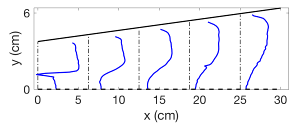 Poor flow development (Inefficient momentum exchange)

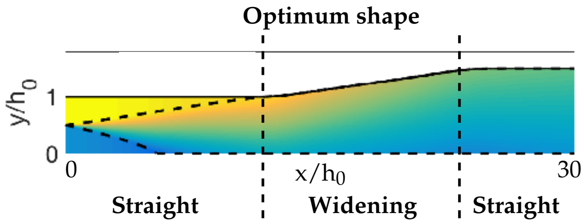

Figure 5 – Table illustrating the advantages and disadvantages of shallow and wide expansion angles when considering the optimum shape of the channel (experimentally measured velocity profiles are used as an argument for flow develoment). Below is illustrated the optimum shape, as well as a colour plot of the velocity and an indication of the position of the shear layers.

3 Decision variables

The main thing we can alter is the shape of the mixer/diffuser. Since the shape is given by a continuous function (e.g. a curve on a plot) rather than a single number (e.g. a point on a plot), this makes the optimisation trickier. However, we can solve it exactly in some specific cases using theoretical tools from *Optimal control theory*. Or, in more general cases, we can solve it by breaking the curve into a finite (but large) number of points, and treating each of these as a decision variable.

Other decision variables include the following:

- The velocity ratio between the secondary and primary flows $\beta = U_2/U_1$.
- The flow rate fraction between the secondary and primary flows $\mathcal{F} = Q_2/Q_1$.

3 Constraints

There are only a few rules to consider. In practical terms, the geometry may be constrained to fit within a desired region, due to either environmental, legal or financial costs. However, for our purposes, we ignore such constraints to keep the problem as simple as possible. Secondly, we require that our optimisation does not deviate from physical reality. Therefore, whilst we optimise our decision variables, we ensure that our mathematical model is being solved to a high degree of accuracy. Since we have validated our model against experimental data, this is an acceptable approach.

3 Results

In Figure 5 we show the typical form of the optimum shape. There are three main insights that we took from the optimisation:

1. The shape is of the form: straight, widening, straight[1].

[1]Note that in some extreme cases of small velocity ratio β, an initial contracting region is optimal. In this situation, the shape is of the form: contracting, widening, straight.

2. In the widening part, the expansion angle is a perfect balance between the effects of mixing and friction.
3. The effects of boundary layer separation are irrelevant.

Firstly, as can be seen in the Figure, the shape is naturally divided into three sections: an initial straight section, followed by a widening section of near-constant expansion angle, and then a final straight section. The initial straight section allows the shear layers to grow and the flow to develop before entering the widening section. Once the flow is sufficiently developed, the widening section reduces the wall drag. Finally the straight section near the end allows any further flow development to occur.

The expansion angle in the middle widening section corresponds to an interesting balance between the physical effects of drag and friction. As we illustrate in Figure 5, shallow expansion angles tend to allow the non-uniform flow to develop well, such that the secondary and primary flows exchange momentum efficiently. However, the large drag due to the narrow walls of the channel results in large energy loss. On the other hand, wide expansion angles result in greatly reduced drag, but also cause an accentuation of the non-uniform flow, where that the primary and secondary flows do not mix well, and lots of energy is dissipated. We discovered that the optimum expansion angle is the compromise between these two effects, and is given approximately by the formula

$$\alpha \approx \tan^{-1}\left(\frac{1}{2}\left(3fS_c^2\right)^{1/3}\right), \tag{6}$$

where f is the *Darcy friction factor* and S_c is the *spreading parameter* which characterises the growth rate of the shear layers. Experiments were used to determine the values of these parameters, such that $f \approx 0.01$ and $S_c \approx 0.18$. Using these values, we find the optimum is a very shallow expansion angle of around 2.8° (total angle 5.6°).

Finally, the optimum expansion angle is considerably smaller than the typical angle associated with boundary layer separation. In the conventional design of flow diffusers, where the flow is uniform and there are no shear layers, the expansion angle is as wide as possible without causing the boundary layers separate. For cylindrical diffusers, this is typically around 3.5°, which is larger than we have deemed optimum for the non-uniform case. Hence, it is clear that in our case, the effects of mixing trump the effects of boundary layer separation, revealing that non-uniform flow diffusers must be treated differently from conventional uniform flow diffusers.

We derived a new expression for the optimum expansion angle of the diffuser, which evaluates to $\alpha \approx 2.8°$.

4 The addition of swirl

In the shape optimisation problem above, the optimal diffuser shape had an initial straight narrow section which was necessary for mixing the flows before expansion. However, this narrow section has the strongest wall drag of the entire geometry. Therefore, if it were possible to mix the flows over a shorter length scale then the narrow section could be made shorter, thereby reducing pressure losses due to drag.

In many engineering applications, there is a mechanism used to aid the mixing of two regions of flow: the addition of *swirl* - that is causing the flow to rotate about its axis of motion. One example is mixing fuels in a combustion chamber. Therefore, as a final avenue of research, we investigated the effect of adding swirl to the secondary flow as a means of accelerating the exchange of momentum with the primary flow. From a practical point of view, this would be easy to achieve, since the turbine could provide a source of swirl naturally.

Our investigation found that, whilst swirl does indeed accelerate the mixing between the two flows, it also dissipates a lot of energy in doing so. We found that the the overall pressure recovery coefficient (5) decreases with the addition of swirl, indicating that the best outcome is achieved by inputting no swirl at all.

5 Case study and recommendations

As discussed earlier, the goal of Venturi-enhanced hydropower is to strike a balance between maximising power and minimising costs. We formulate a multi-objective optimisation problem for both the cost and power. The power is linearly related to the

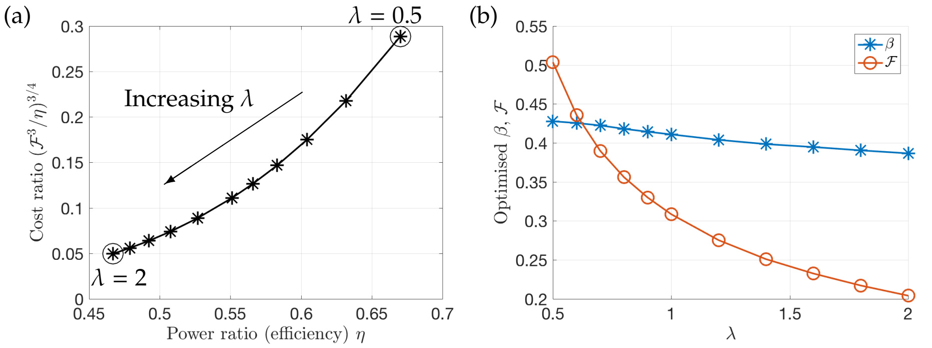

Figure 6 – Multi-objective optimisation of power and cost. (a) Pareto distribution obtained by varying $0.5 \leq \lambda \leq 2$. (b) Velocity ratio β and flow rate fraction $\mathcal{F}$, optimised for different values of the multi-objective parameter λ. Large λ corresponds to prioritising cost.

efficiency η, given by (3), whereas we prove that the cost is linearly related to the quantity $(\mathcal{F}^3/\eta)^{3/4}$, involving both the flow rate fraction and the efficiency. Therefore, we choose the combined objective:

$$\text{Maximise}: \quad \eta - \lambda \left(\frac{\mathcal{F}^3}{\eta}\right)^{3/4}, \tag{7}$$

where $\lambda \geq 0$ is a control parameter which determines the relative importance of each objective. The parameter λ can also be related to the expected return on the investment (which is related to the efficiency) over the lifetime of the project.

A multi-objective optimisation is where two or more different quantities are optimised (such as cost and power), where there is a given weighting between these two objectives.

5 Case study

As a case study, we choose a hypothetical river/dam with a total head drop of 3 m that provides a flow rate of 4 $\mathrm{m^3\,s^{-1}}$ to the Venturi-enhanced hydropower. The multi-objective optimisation parameter λ is varied between 0.5 and 2 to illustrate the effect of altering the relative importance of cost and power (large λ corresponds to prioritising cost). In each case, we optimise not only the channel shape, but also the other decision variables, the velocity ratio β and the flow rate fraction $\mathcal{F}$.

We plot the distribution of outcomes, known as the *Pareto distribution*, in Figure 6 (a). The optimum power ratio (or efficiency) varies between 0.47 and 0.67 (corresponding to a total output power of 55 kW and 79 kW), whilst the cost (or turbine volume) ratio varies between 0.29 and 0.05. In the most expensive limit, the Venturi-enhanced hydropower turbine is 29% of the cost of a conventional hydropower turbine and generates 67% of the power available in the water[1]. In the cheapest limit, it is 5% of the cost, whilst generating 47% of the available power.

In Figure 6 (b) we plot the optimised values of β and $\mathcal{F}$ as we vary λ. As expected, increasing λ, and hence the importance of cost, results in a lower value of the flow rate fraction. In the most expensive limit the turbine receives 50% of the total flow rate, compared to 20% in the cheapest limit. The optimum value of β does not vary much with λ, but decreases slightly as λ increases. Hence, for all cases we can take $\beta = 0.4$ as a good estimate for the optimum value. The optimum channel shape is very similar for each of these cases, and is close to the shape illustrated in Figure 5.

5 Recommendations

The current design of VerdErg's hydropower geometry is actually already very similar to the optimum shape we identified. That is to say, their channel shape consists of a straight section, followed by a widening section, followed by a straight section (though our optimisation output was entirely independent of this). Therefore, the main result of our optimisation is to prove that no higher efficiency can be achieved than what is there already. Our optimisation has given useful insight into the explanation behind the optimum shape, including the interesting balance between the effects of mixing and

[1]Note that in the case where the conventional hydropower turbine is 100% efficient, then the power generated by the conventional hydropower turbine is equal to the power available in the water.

friction, and the fact that the optimum angle of around 2.8° is below what is typically used for conventional diffuser design.

In our case study example, we have shown that by prioritising either cost or power, a variety of different outcomes can be achieved. Depending on the economic model chosen by VerdEerg, a suitable configuration can be chosen based on our recommendation. The current design of VerdErg corresponds to a flow rate fraction of $\mathcal{F} = 0.2$, and the efficiency has been observed to be around 50%. We have shown how the efficiency could be increased to around 67% if a flow rate fraction of $\mathcal{F} = 0.5$ were chosen. However, as we have shown, this would cost around 6 times as much.

Our final recommendation is to avoid swirl at all costs. VerdErg were interested in using swirl as a mechanism to improve efficiency, but we have shown that swirl can only provide a negative effect. In addition, since the turbine may already be producing a small amount of unintentional swirl, efforts should be made to remove this.

6 Discussion and conclusions

We have studied the fluid dynamics of Venturi-enhanced hydropower to predict and optimise the efficiency. We have identified the important phenomena that play a role in how the primary and secondary flows mix together, and created a mathematical model which serves as a useful predictive tool. We used this mathematical model, in conjunction with optimisation tools, to find the optimum channel shape. This revealed an interesting balance between the effects of mixing and friction, and an expansion angle below the conventional design of diffusers. We showed that swirl in the flow may help to accelerate mixing, but causes a loss of energy in doing so, and should therefore be avoided. Finally, a case study reveals a distribution of outcomes that can be achieved based on the priority of capital costs and hydrodynamic efficiency.

The project with VerdErg has not only shed light on the design of Venturi-enhanced hydropower, but has also opened doors into some deeper and fundamental questions in the field of fluid dynamics. Such questions include how shear layers grow when confined in a channel, and if there is an optimum expansion angle that balances the effects of mixing and friction.

7 Potential impact

Our mathematical model for the mixing of the primary and secondary flows will serve as a predictive tool for the design of Venturi-enhanced hydropower. For a given shape, our model can predict the efficiency, as well as the sensitivity of the design to changes. Our optimisation results can be used to ensure that the design is at the best possible configuration. And finally, our case study illustrates how our optimisation tools can be used to explore a tradeoff between prioritising power and cost.

Dr Paul Bird, chief engineer at VerdErg Renewable Energy said: "*While the physical realisation of the VETT hydropower generator is quite simple (important in keeping cost down) the fluid dynamics processes are anything but simple. We started with a good understanding of the basic energy flows and their relationships to geometry, but optimising the flows to reduce losses to a minimum turned out to be difficult. Our main development approach was to carry out physical modelling at quite a large scale (500 litres per second) supplemented by computational fluid dynamics studies. Each design iteration gave a reduction in the losses, but we felt there were still more improvements to be made. The work described here by Graham Benham added valuable insights into the energy exchange processes, and the relative magnitudes of different losses, together with a number of optimisation tools. This work will assist us in making VETT fulfil its potential to make a valuable contribution to generating clean renewable energy worldwide.*"