The 2nd Generation Turbine Lander: Design, Analysis, and Testing

Christopher Bassett, Harlin Wood, Paul Gibbs, Ben Cunningham, Jesse Dosher, and Tracy Tran

Technical Report
APL-UW TR 2503
August 2025

Applied Physics Laboratory
1013 NE 40th Street

University of Washington
Seattle, Washington 98105-6698

ACKNOWLEDGEMENTS

This report was sponsored by the Office of Naval Research, Code 321US, Undersea Signal Processing. Funding was provided through Naval Sea Systems Command contract N00024-21-D-6400 under Task Order N00024-22-F-8714.

We acknowledge the contributions from faculty, staff, and graduate students including Brian Polagye, Abigail Snortland, Ari Athair, Aidan Hunt, Greg Talpey, and Owen Williams. This group participated in critical discussions to identify new design parameters within the project constraints and provided feedback on predicted individual blade loads. This feedback was derived directly from prior laboratory-scale experiments; constraining options for the redesign effort would have been considerably more difficult without their input. We also thank APL-UW engineer Justin Burnett for participating in several mechanical design reviews.

CONTENTS

LIST OF TABLES

LIST OF FIGURES

EXECUTIVE SUMMARY

The 1st generation Turbine Lander, a small-scale (1 m^2) vertical-axis cross-flow turbine on a gravity foundation was deployed in Sequim Bay, Washington for 141 days from October 2023 to March 2024. In the 1st generation system, which represented a laboratory to field effort, a priority was placed on survivability over efficiency and, when possible, leveraged commercial off-the-shelf components instead of custom component development. Analysis of pre-deployment characterization data and performance data acquired in situ reveals a broad range of system inefficiencies and design modifications that could improve performance. At the same time, efforts to operationalize the concept following demonstration of feasibility required improvements to the system to enable autonomy. Specifically, the 1st generation system required 3-phase 480 Vac power input and included no integrated energy storage. The power electronics, designed for industrial automation, facilitated the acquisition of high-quality data, but at the cost of unsustainable high hotel loads for long-term deployments.

This project sought to address fundamental limitations of the 1st generation system design to advance the concept towards a hypothetical autonomous deployment in moderately energetic environments (i.e., maximum current speeds < 2.5 m/s). Based on lessons learned the following areas of focus emerged as design priorities for the 2nd generation system:

- Increase power capture of the rotor by modestly increasing its size, modifying its aspect ratio, and adjusting the preset angle of attack to enhance torque generation
- Decrease the potential for biofouling by minimizing protruding components (e.g., fasteners) and removing crevices in which floating flora could become lodged
- Minimize parasitic mechanical losses associated with seals and lubricants
- Leverage design modifications to enable the implementation of a broader range of control strategies, particularly those that can operate without inflow current measurements
- Identify and implement an alternative power electronics package to reduce hotel loads and generate direct current power to facilitate integration with energy storage

These objectives have been achieved and this report summarizes the design, fabrication, and characterization efforts involved with the development of the 2nd generation Turbine Lander system. First, we summarize the 1st generation system and previously recommended approaches for improving the system. The new power electronics, which were integrated in collaboration with researchers from the Monterey Bay Aquarium Research Institute (MBARI) are described. A new battery system, which is being developed in parallel to this project under a different contract, is introduced. Next, the new rotor design is introduced and supporting engineering design efforts, both analytical and numerical, are summarized. The report concludes with detailed performance characterization of the new rotor and power take-off with comparisons to the 1st generation unit and recommendations for future work.

These efforts have resulted in significant performance improvements. For example, the 1st generation system could not generate power until currents exceeded 1 m/s, and that value excluded the hotel loads of the power electronics. The new system achieves net power generation at inflow speeds less than 1 m/s after accounting for internal power consumption by the power electronics. As inflow speeds increase more significant gains in power generation have been achieved with the new design producing more than 2.3 kW at 2.5 m/s, over 800 W greater than the initial design. At inflow speeds between 1.25 and 1.75 m/s increased power outputs of greater than 200 W were achieved. These differences represent the integrated improvements achieved not by a single,

significant modification to the system, but a large number of incremental changes and provide a pathway for achieving performance improvements in other small-scale systems where differences as small as a few 10's of Watts determine concept feasibility.

1 INTRODUCTION

The 1st generation Turbine Lander (Figure 1), a small-scale, vertical-axis, cross-flow tidal turbine, was deployed in Sequim Bay, Washington for its first in situ test in October 2023 [2]. During its 141-day deployment the Turbine Lander successfully generated power during periods with strong currents until operations were curtailed in mid-February following the failure of a second blade while the system was operating and generating power. Before operations were curtailed the unit had been active in strong currents for more than 1000 hours, over 900 of which occurred during the deployment in Sequim Bay. The failure of the foils in situ was ultimately attributed to a combination of collisions with large objects in the water, likely floating logs that had been drawn into the water by perigean spring tides that corresponded to both blade failure events, and a design flaw in the assembly that retained the pins that attached the foils to the rotor.

While the failure of the foils represented the most significant problems encountered during operations of the 1st generation rotor, prior characterization following system fabrication identified several pathways for improved system efficiency. Furthermore, optical cameras used to record environmental interactions [3] identified issues with the rotor design that allowed eelgrass to regularly foul the foils and the shaft, leading to operational inefficiencies.

The identified pathways likely to lead to notable improvements in power generation included:

- Minimization of mechanical losses that represented parasitic torques
- Reducing the drag profile of exposed and rotating components
- Biofouling mitigation measures
- Modifications to the rotor's size and aspect ratio to increase power capture and torque generation.

In addition, the 1st generation unit requires 480 Vac power provided by a shore cable, even if the system is producing power, which is dissipated by a resistive load dump. Moving towards an autonomous system would require modifying the system's power electronics and the development of integrated energy storage to enable system autonomy.

This report summarizes modifications to the system, which we now consider the 2nd generation Turbine Lander. The following sections describe modifications to the system's power electronics, development of a modular battery system, and activities related to the redesign, fabrication, and testing of the rotor. The 2nd generation system has not been deployed in the field, but vessel-based testing on R/V *Russell Davis Light* (*RDL*) are compared to the 1st generation system.

1.1 First-generation Design and Identified Inefficiencies

To understand the system's modifications we first review the driveline and power take-off (PTO) of the 1st generation Turbine Lander. This work is described in greater detail, including part numbers and dimensions, in Bassett et al. [2] and its associated supplemental material. The vertical-axis cantilevered rotor sits above the PTO, which is housed in the unit's foundation (Figure 2). This summary begins with the rotor and works down to the generator, highlighting features of relevance to the redesign effort.

The 1st generation rotor was 1.19 m tall and 0.85 m in diameter and consisted of four, straight composite foils with NACA 0018 profiles. The foils were fabricated such that they had a 9° pre-set

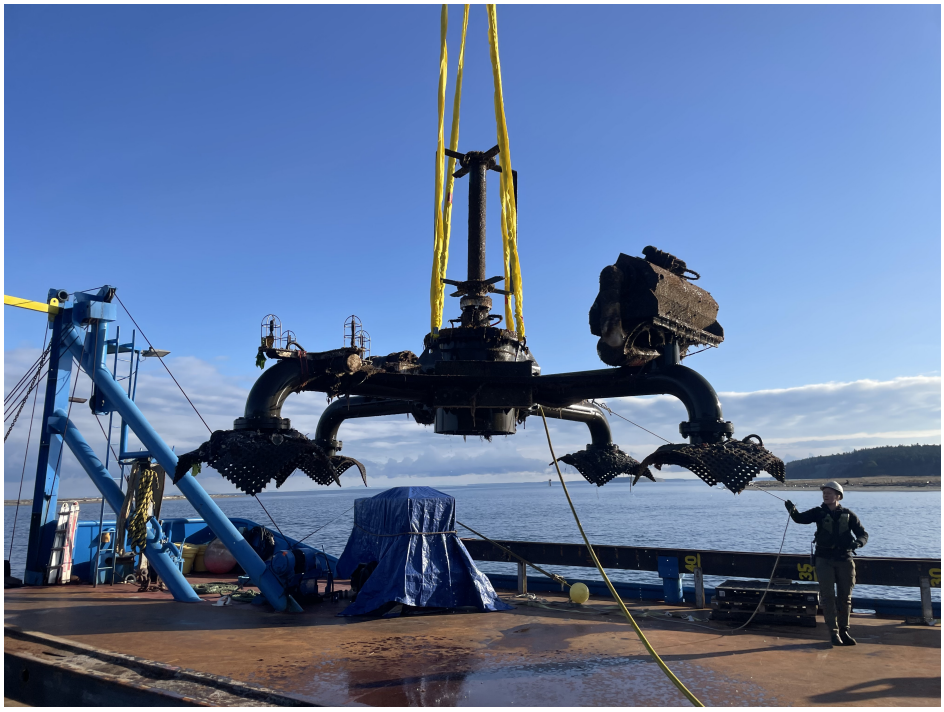

Figure 1: The 1st generation Turbine Lander recovery from Sequim Bay.

pitch angle relative. This angle was larger than intended and a result of a design error, which resulted in operational inefficiencies. The struts that connect the foils to the hub/endplates and shaft were 1.27 cm (1/2") thick with convex fillets with radii that were one-half of the thickness, resulting in a rounded edge. The large (15.24 cm [6"]) rotor shaft is a result of predicted loads and efforts to limit the deformation of the shaft under peak loads and cannot be significantly modified without cascading design modifications. However, while the portions of the rotor further from the shaft were generally streamlined, the endplates securing the shaft to the rest of the assembly had protruding fasteners that were exposed to the flow. These were ultimately a source of drag, albeit small due to their size and proximity to the axis of rotation, and an occasional source of modest biofouling.

At the base of the rotor the unit was directly coupled with the driveshaft and bearing pack. Due to height limitations, accommodating the peak loads during operation required that the bearing pack consist of two stainless steel tapered roller bearings. To minimize the risk of corrosion and fouling of the bearings due to exposure to particulate laden saltwater, the bearing pack was isolated using two lip seals and an exclusionary v-seal. The bearing pack was filled with a biodegradable oil and maintained above ambient pressure using a pressure compensator. At the base of the bearing pack, a magnetic coupler installed on the generator housing joined the rotatory motion of the rotor shaft into the housing without the need for a dynamic seal. Laboratory-based analysis identified that the lip seals each contributed about 5 N-m of parasitic torque, resulting in combined losses of between 60 to 120 W across the rotational speeds of the system while generating power. Viscous losses associated with the oil-filled bearing pack could not be measured, but we estimate them to be between 10 and 30 W during operations based on analytical models combining the geometry, relative speed of the shaft, and oil viscosity.

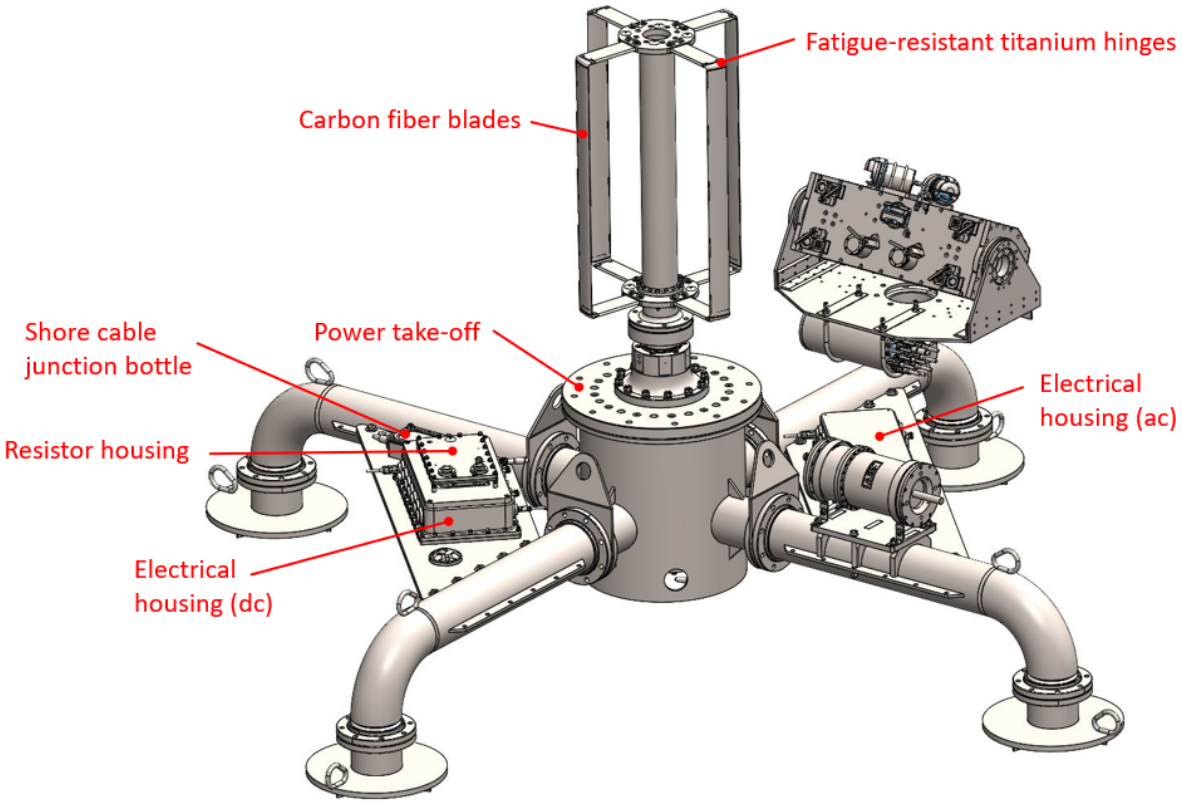

Figure 2: An annotated render of the Turbine Lander equipped with the 1st generation rotor and electrical housings. The footprint of the foundation is approximately 4 m by 4 m and the total height is approximately 2.7 m.

The 1st generation system generator was located in an atmospheric pressure (air- or nitrogen-filled) housing that contains nothing but magnetic coupling, shaft, generator, and supporting structure, although the initial design included fans for active cooling of the generator that were found to be unnecessary after vessel testing. As installed, the generator, which is used primarily in industrial automation applications, included a large diameter shaft seal that had not been removed. Only after the deployment of the 1st generation system were the losses associated with this seal measured directly, at which point they were determined to be comparable to the combination of both lip seals in the bearing pack and a significant portion of total parasitic losses in the system.

Industrial automation electronics controlled the 1st generation system. Given their typical use, and inherent flexibility, these components are large (see Section 2 for details) and not optimized for applications in which minimal hotel loads are required. In addition, the electronics required an internal control computer and supporting equipment leading to the inclusion of additional components to support communications and system stability, particularly in cases where shore power was lost.

Review of the full driveline/PTO assembly yielded several opportunities for improvement in the 1st generation system. First, the rotor could be improved to mitigate biofouling, improve hydrodynamic inefficiencies, and address prior failure points. In addition, if consistent with other design requirements, the rotor geometry could be modified to increase power capture by increasing

its size or modifying its aspect ratio. By selecting a new, lower viscosity oil for the bearing pack modest power gains across operating states could be achieved. The removal of any seals not critical to the survivability of the system would yield gains. Lastly, to facilitate autonomous deployment and minimize hotel loads, new power electronics were required. These factors, identified through detailed laboratory and field testing, motivated the system modifications for the 2nd generation system.

Figure 3: The 1st generator rotor installed on the Turbine Lander prior to deployment in October 2023.

2 POWER ELECTRONICS

The 1st generation system was designed around a Siemens motor and Siemens industrial automation power electronics, which are described in greater detail in Bassett et al. [2] Given that this equipment is found typically in industrial and commercial environments, it is relatively large. Furthermore, as a result of a broad range of capabilities designed to support industrial applications, the units are inherently power intensive, at least when compared to targeted levels of power production for a small-scale tidal turbine deployed in coastal environments with moderately energetic currents. While not a logical fit for long-term development, Bassett et al. [2] notes that this equipment was selected for the 1st generation system due to their inherent flexibility, support for data acquisition at relatively high data rates, and a lack of alternative commercial off-the-shelf components suitable for the project.

The 1st generation power electronics had two housings (Figures 4 and 5); one received 3-phase 480 Vac from the shore-side cable and a second contained direct current components including a motor module, computer, relays, and a supercapacitor. Rectangular housings, both 60 cm x 40 cm x 30 cm or larger, were selected to limit the size and excess buoyancy that would result from extra space in the housings. To dissipate power generated by the system, an external housing was equipped with three resistors capable of dissipating 1.8 kW of electrical power.

Three primary issues motivated the adoption of new power electronics: size, hotel loads, and

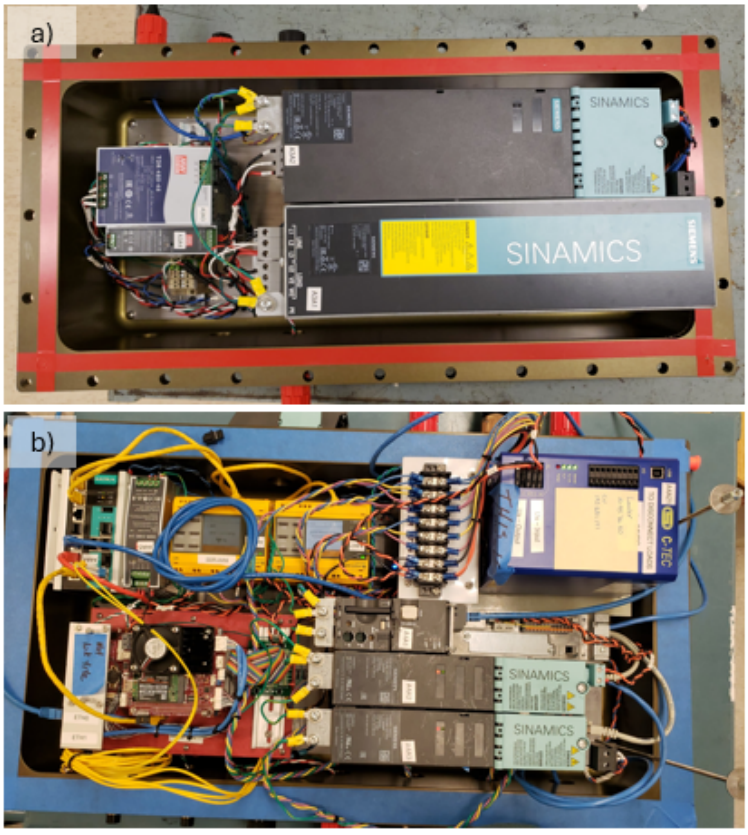

Figure 4: The two power electronics housings, a) ac and b) dc, from the 1st generation Turbine Lander.

the development of energy storage requirements, which motivated the shift to dc power. Combined with improved rotor performance, this shift in operations would facilitate operations without a cable to shore and allow future deployments to be performed at a broader range of sites, including those that lack easy access to shore-side facilities and 3-phase 480 Vac power. Furthermore, after accounting for non-recurring engineering costs in development, a custom system would, ideally, decrease costs.

2.1 MBARI Power Converter

We reviewed commercial off-the-shelf (COTS) systems and did not find any obvious solutions. We turned to Monterey Bay Aquarium and Research Institute (MBARI), which developed its own wave energy converter (WEC),[4] the MBARI-WEC. Similar to our experiences, MBARI did not identify suitable COTS components for use with their high-speed hydraulic motor and therefore decided to develop their own custom unit. Although the Turbine Lander's direct-drive generator is suitable for low-speed, high-torque operations, the motor's torque and voltage constants coupled with the rotor's targeted operations turned out to match well with the custom MBARI WEC power converter, referred to here as the MPC. Thus, we chose to further evaluate the potential to use the MPC in the Turbine Lander. At a high level, the MPC rectifies the three-phase voltage and boost it to a constant voltage (325 Vdc by default) suitable for charging batteries. The system is also designed to integrate with a load dump and can accommodate bidirectional power support complex controls schemes. The system requires battery storage to supply the modest hotel loads when not generating power, but is self-powered if sufficient power is being generated (this contrasts with the 1st generation Turbine Lander system that relies on external power under all circumstances).

In its modified configuration the MBARI power electronics report several system variables via Controller Area Network (CAN) at up to 100 Hz during operations. These include winding current and window current setpoint, shaft speed and shaft speed setpoint, winding voltage, targeting winding voltages and currents (as determined by a control loop), currents to the load dump, currents to the battery, and PID gain settings. System control loops can be modified as necessary, but the default is a PID loop established to maintain a targeted control torque (or winding current).

Evaluation of the MPC took place is several steps. First, modest modifications to the Lander generator were made and the PTO was installed on APL-UW's dynamometer for proof of concept testing, in which the Siemens gear was swapped for a spare benchtop MPC. This required first replacing the generator's absolute encoder with an off-the-shelf incremental encoder sourced from Encoder Products (part number: 260-N-R-10-L-6000-R-HV-1-S-SF-2-N, line count 6000) that could meet the requirements of the MPC. Once in hand, a new mechanical mount was to join the new encoder to the generator enclosure. The test configuration is described in Figure 6. No investments in design modifications to either system were made in advance, so the MPC was seated on a desktop and wired to a laptop computer, dc power supply, and a load dump in an ice bath. The prime mover for the dynamometer system, which is described in Bassett et al.[2] was the same model Siemens motor (1FW3202) and a torque cell installed in-line between the prime mover and PTO measured the input shaft speed and torque.

Testing was carried out under speed control (i.e., the prime mover was controlled to maintain a constant shaft speed while commands provided to the MPC controlled the currents in the windings) and sampled all four operational quadrants, meaning that the system was operated with both positive/negative shaft rotation speeds and positive/negative control torques. Shaft speeds of –100

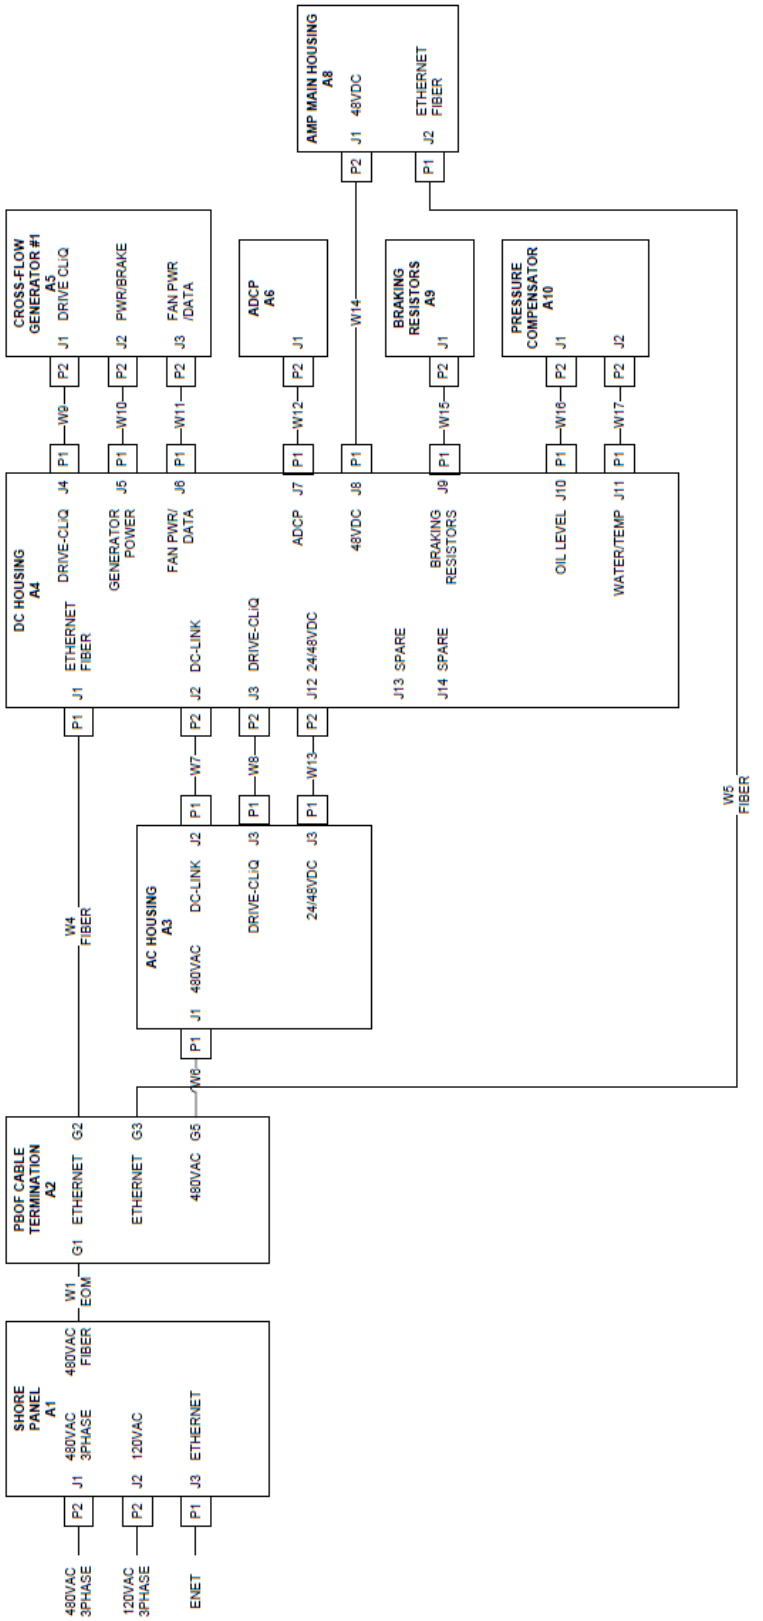

Figure 5: A block diagram of housings and critical electrical components in the 1st generation Turbine Lander.

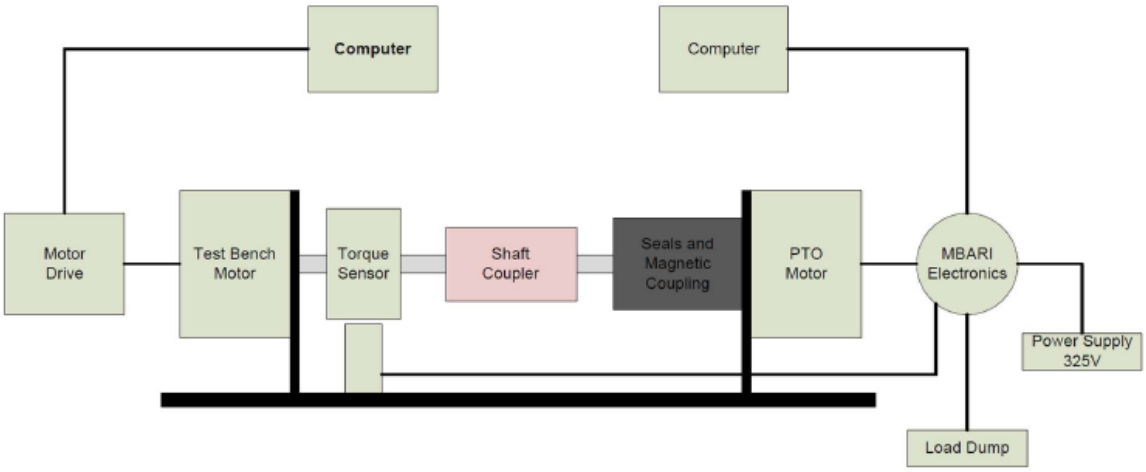

Figure 6: The benchtop testing configuration for the MPC with the Turbine Lander PTO.

– 100 RPM (including 0 RPM) were sampled in increments of 10 RPM while winding currents of –20 – 20 A, in increments of 2 A, were tested. Given the motor's torque constant (25 N-m/A), winding currents in excess of approximately 17 A, depending on the signs of the shaft speed and current, can exceed the PTO's torque limits that are driven by the magnetic coupling. The upper limits of the test conditions exceed the torque that can be generated by the Turbine Lander rotor's and therefore represent an operational envelope than can be achieved by the system during deployment. While only the power generating quadrants are of general interest for power production, the broader parameter space allowed for mechanical losses to be inferred and was also used to verify that the MPC behaved well with the Lander PTO across its potential operating space.

In testing we verified that the MPC was able to meet all of basic requirements that would be required to operate with the Lander PTO and that we could take next steps to develop and integrate the system with the Lander. In addition, we were able to calculate the sources of electrical losses from the system under different operating states (Figure 7). To effectively map system losses the system resistance values were measured across critical components and used to map I^2R, where I is electrical current and R is resistance, and combined with the current values measured by the MPC. These calculations were carried out at a shaft speed of 0 RPM, when no mechanical work was being performed by the system. The total electrical losses represent the various restive losses in addition to the transistor losses. From this we found that the transistor switching losses and associated components are approximately 68 W (averaged across the different current settings) with motor winding I^2R losses only exceeding switching losses for control torques greater than 6 A. Under most conditions in which the Turbine Lander is designed to operate (rotor mechanical power $<$ 2 kW) the total electrical losses are less than approximately 200 W (note that this excludes additional losses associated with boosting the output voltage of the generator).

As part of the battery design process (Section 3) alternative target voltages for the power electronics were considered. While the default operates at a 325 Vdc, the target voltage can be modulated to interface with energy storage and is capable of operating over a broad range of voltages. The Lander's generator has a voltage constant of ≈1.5 V/RPM. Over targeted operational ranges for the 1st and 2nd generation systems, the minimum and maximum expected voltages (peak-to-peak) of the generator are 64 V and 200 V (peak-to-peak), respectively. The MPC rectifies and steps up the output voltage of the generator to the target voltage for the system at its operating

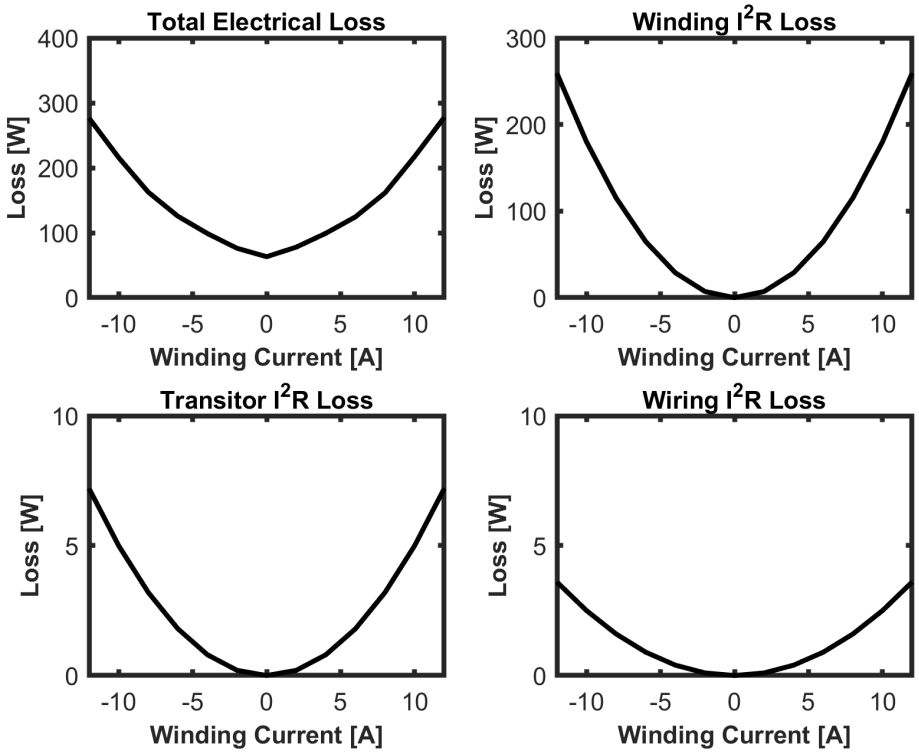

Figure 7: Electrical (I^2R) losses associated with the MPC determined during dynamometer testing at the default voltage of 325 Vdc.

state. The process of stepping up this voltage, especially with large voltage differences, is a source of significant losses.

Though a simplification of the MPC's six-switch power factor correction rectifier, a single switch boost converter provides a framework in support of a lower operational voltage. Stepping up the voltage requires higher duty cycles with a fixed switching frequency. This, in turn, leads to higher conductive losses, lowering overall efficiency. Figure 8 shows example calculations for generator speeds, voltages, and corresponding inflows for two MPC target voltages with this simplified approach. The estimated duty cycle required to step up the voltage is defined by $V_{bus} = V_{gen}/(1-D)$, where V_{bus} is the MPC voltage, V_{gen} is the generator voltage, and D is the duty cycle. Duty cycles greater than 0.6 are expected to lead to significant (non-linear) drops in efficiency. While unavoidable over parts of the turbine's operating range, dropping the bus voltage can significantly reduce the duty cycle thereby improving overall efficiency. Because lowering the bus voltage has the most significant impacts on efficiency when the generator voltages are lowest, this approach can help minimize system losses when the total net power available from inflow conditions is the lowest and generating meaningful power is most difficult.

Examples of data obtained during testing of the 2nd generation rotor with the MPC at different bus voltages is described in Section 5.

2.2 Integration and Dynamometer Testing

Once proof of concept had been demonstrated, a new approach to packaging the MPC had to be identified. As integrated with the MBARI WEC, the MPC is a stack of three printed circuit boards (PCBs), each about six inches in diameter. While this packaging is conveniently sized for most applications, this geometry precluded direct integration with the Turbine Lander. Specifically, our goal was to house the MPC within the generator housing. The 1st generation Turbine Lander had included heat sinks and fans integrated into the generator housing. These ultimately proved unnecessary and they were removed prior to deployment in Sequim Bay. This left a few inches in space at the end of the generator housing, except in the center where the new encoder sits, which was more than sufficient for the MPC.

Direct integration of the power electronics into the generator housing required three significant modifications to our systems. First, a new end cap was machined that would allow the MPC to be mounted directly to the housing. Next, given height restrictions, the PCBs needed to be laid flat in the limited space in the housing. Lastly, an alternative load dump needed to be integrated with the housing. The end product of these three design modifications in shown in Figure 9. Modifications to the interior of the PTO housing endcap were made to allow the boards to be directly secured while the fabrication of a custom flex circuit (by San Francisco Circuits) allowed the board stack to be separated into three components distributed around the edges of the endcap, leaving clearance for the encoder. The load dump housing, which contains six 1500 W water heater elements submerged in oil, is mounted directly to the last portion of the end. The new housing endcap was machined with four through holes for connectorization to support power, communications, and ancillary equipment. In practice, these will support connections to the generator, a battery system, communications, and external inputs (e.g., data from sensors to support controls). These holes are oversized for most standard connections and machined Polyether Ether Ketone (PEEK) washers are used to isolate the connectors and avoid mixed-metal contact. This approach fundamentally reduces the number of independently "housed" components from five (as shown in the right-hand side of Figure 6) to two: The PTO housing and the power supply or battery.

After integration, the modified PTO with the integrated MPC was again installed on UW's

Motor Speed	Motor Voltage	Inflow Speed	Bus voltage [Vdc]		
[RPM]	[Vrms]	[m/s]	240	320	
30	45.60	0.86	0.81	0.86	Duty Cycle
40	60.80	1.15	0.75	0.81	
50	76.00	1.44	0.68	0.76	
60	91.20	1.72	0.62	0.72	
70	106.40	2.01	0.56	0.67	
80	121.60	2.30	0.49	0.62	
90	136.80	2.59	0.43	0.57	

Figure 8: The predicted duty cycles (color) for single switch boost converter with red, orange, yellow, and green corresponding to duty cycles likely to lead to lower losses, respectively. This simple approach does not represent the full MPC and is simply intended to convey the general benefits of reducing the bus voltage. Note the rms voltage and inflow speeds are calculated directly from the speed of the generator, its voltage constant (1520 V_{rms}/1000 RPM), and an assumed tip speed ratio of 1.9.

dynamometer. Having previously confirmed the suitability of the system, the second round of testing sought to confirm the new PCB layout with the flex circuit operated as expected and to develop modified PID controls settings to account for the differences between the operations of the MBARI WEC and the Turbine Lander. In the case of control settings, the approach was to best simulate hydrodynamic torques applied by in-situ or vessel testing to minimize the amount of vessel time required to make the system operational. Thus, three specific sets of tests were performed:

- MPC performed constant speed control with the prime mover stepping through constant torque settings
- MPC performed constant torque control with the prime mover stepping through constant speed settings
- MPC performed speed control with the prime mover applying oscillatory loads representing phase-resolved PTO control torques from first-generation characterization

In addition to demonstrating stability of the controls, this would also allow us to generate (and dump) power as would be expected in situ.

Successful testing on the dynamometer required that the MPC maintain generator speed at a relative constant while changes in torque were applied by the prime mover, and that constant control torques could be maintained while the prime mover altered shaft speeds. The most complex laboratory test carried out prior to full system field testing was speed control under oscillatory loads representing typical rotor-driven torque fluctuations. An example of speed and torque data from the prime mover and MPC are shown in Figure 10. Similar data obtained across the operational range generated the confidence required to transition the system to R/V *Russell Davis Light* for further testing. Further modifications to control gains were ultimately required (as expected) given that the dynamometer could apply arbitrary torque commands, but the moment of inertia and added mass associated with the rotor were not captured in the laboratory tests.

Figure 9: MPC on the PTO housing endcap with the load dump (near the bottom).

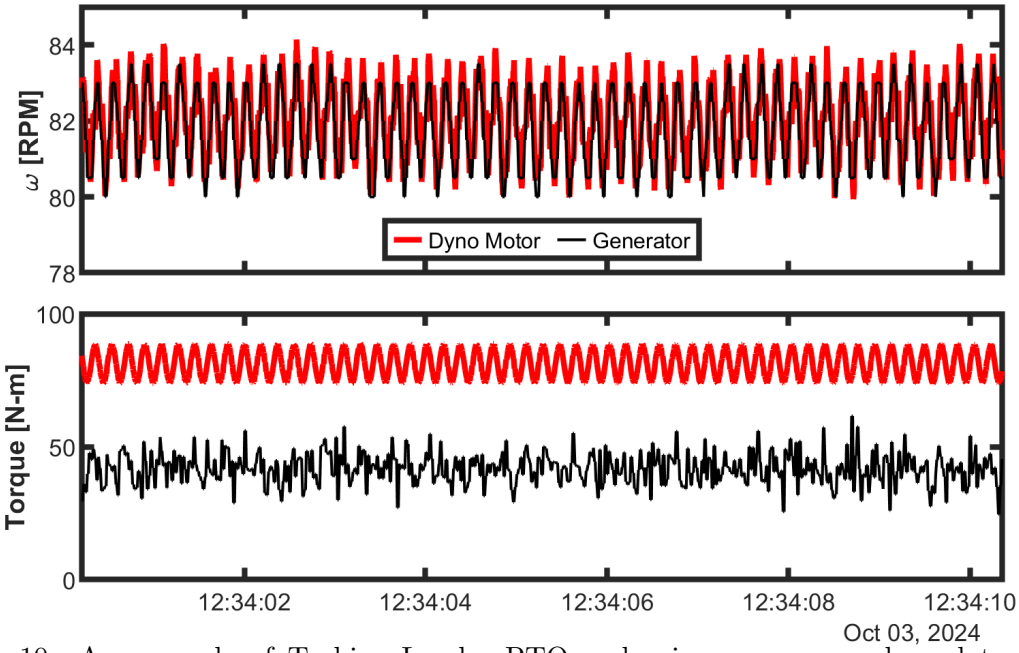

Figure 10: An example of Turbine Lander PTO and prime mover speeds and torques during oscillatory load testing. The prime mover was programmed to oscillate between applied torques of approximately 74 – 88 N-m while the PTO was set to maintain a constant speed of 82 RPM. This control was achieved with oscillations of ± 2 RPM.

2.3 Representative Field Testing

The transition to *RDL* for representative field tests using the 1st generation rotor initially produced surprising, and frustrating, results. While adequate controls had been demonstrated in laboratory testing, initial attempts to implement speed control algorithms in situ were unsuccessful and the implementation of constant torque control approaches was restricted by the performance of the initial rotor. The fundamental challenges inherent in this early testing were related to gain settings programmed into the MPC. In its initial design application the unit operates on a wave energy converter and is therefore responsive to inputs with large dynamic ranges that occur over periods more than an order of magnitude longer than the oscillatory torques produced by the turbine. Basic control features of the system were, at the time, stored on the control board in the housing. A new user interface to allow remote control of the system was set up, at which point additional testing was carried out with the goal of identifying a set of PID control gains that would result in relative stability in constant torque and speed control modes.

Initial testing was carried out with *RDL* at a constant speed of 1.5 m/s. Because appropriate control gains were unknown, the system was started at an RPM near the peak of the efficiency curve ($\lambda = 1.9$) and the variability in control torques and RPM were monitored over time. First, proportional control gains were established that supported adequate transitions to new rotor speed set points when provided. While vaguely defined, we considered a stable transition to a new set point, without significantly overshooting the target, within several seconds to be suitable. Once a suitable proportional gain was identified at a single speed, additional testing and iteration occurred for inflow conditions at 1 – 2.5 m/s. With a suitable proportional gain established, integration gain was added to the control loop to minimize stead-state errors and the measurement noise in

the system. A derivative gain was also implemented, but we ultimately targeted a relatively weak value out of concern that a strong gain would result in oscillations. The control gains identified in this iterative process remained unchanged throughout subsequent efforts to characterize the 2nd generation rotor and implement more complex control algorithms.[a]

2.4 Power Electronics Summary

A comparison between Figures 4 and 7 illustrates the reduction is space and equipment achieved by the adoption of the MPC. Specifically, what can now be housed within the PTO housing previously required two large housings that, combined, weighed more than 45 kg. The integrated load dump removes an additional housing from the system. The MBARI system's primary losses at low power states are the transistor losses, but the microcontroller can enter a low power consumption state (<4 W), functionally shutting down the system while it awaits commands or is not generating power. The reduction in hotel loads due to the more limited operational capabilities represents a significant step forward in developing small-scale autonomous systems that can maintain high persistence by minimizing unnecessary power requirements. While this system represents the next step towards power generation and energy storage for the system, we have maintained the ability to operate with the 1st generation power electronics in an application where three-phase 480 Vac is desired. This can be done by swapping the encoders, removing the MPC, and changing the bulkhead connectors. While we cannot evaluate the current costs of the 1st generation power electronics due to a lack of recent quotes, records from the prior build suggest that the costs of the MPC are less than 50% of the costs of the Seimens gear. Including the costs of large housings, the cost-savings associated with the 2nd generation system are significant.

[a]Note that specific gain values determined in the process and later applied in testing are not provided. The MPC's control loop, when developed, used gains that were not specifically coupled to engineering units. Thus, its specific control gains, if mentioned, would not provide clarifying information. For qualitative purposes we note that the proportional gain was set to have a strong response to errors between the desired setpoint and actual operations while the integral gain was set to correspond to periods of several seconds. Integral gains corresponding to this time scale are long enough to capture several rotations yet generally shorter than significant changes in inflow conditions in situ. Rega*RDL*ess, these gains will need to be re-evaluated once the system is redeployed in the field.

3 BATTERY

The development of a battery system is critical to facilitate autonomy of the Turbine Lander given that the 1st generation system was cabled and lacked inherent storage capabilities. Development of a battery system was ultimately dependent on the development of the 2nd generation power electronics. A battery system is currently under development, supported under contract NAVSEA N00024-21-D-6400, Task Order N0002423F8719. It is described in general terms here.

APL-UW worked closely with MBARI to develop design specifications and identify a vendor for the custom battery. Full details are beyond the scope of this work, but the critical specifications are included in Table 1, while Figure 11 includes a block diagram of the system features. The nominal voltage is 270 Vdc, although the system could be reconfigured for different voltages. Its storage capacity, 9.6 kWh, is designed to support loads requiring an average of 100 W for several days if the system goes through a period with little to no power generation (e.g., neaps tides or calm wave conditions). To support both wave and current energy applications, the system can support instantaneous currents up to 35 A. Time-averaged current limits are fundamentally limited by thermal management requirements but will generally exceed time-averaged power production for all but the most energetic conditions. The system can readily accommodate rapid charge/discharge cycles and permits bidirectional flow at short timescales to maximize system efficiency during operations.

The battery operating voltage has implications for its size, storage capabilities, and for the efficiency of power conversion from the marine energy converter with which it interfaces (in this case either a tidal or wave energy converter). MBARI and APL-UW analyzed prior operational data from the MBARI WEC and Turbine Lander to identify both upper limits probability distributions of operational voltages to constrain the pack voltage. Based on these distributions, which commonly operate well below 100 V for both systems, a nominal pack voltage of 270 Vdc was selected. This decision, described in Section 5, has led to a shift in the target voltage range for the MPC in this application. A voltage versus state of charge (SOC) curve for the system is shown in Figure 12. The target range for suitable state of charge is between 240 and 307 Vdc. At states of charge less than 25%, additional operational modes will be built into the system to avoid critically low SOCs that could damage the pack or lead to the shutdown of critical components.

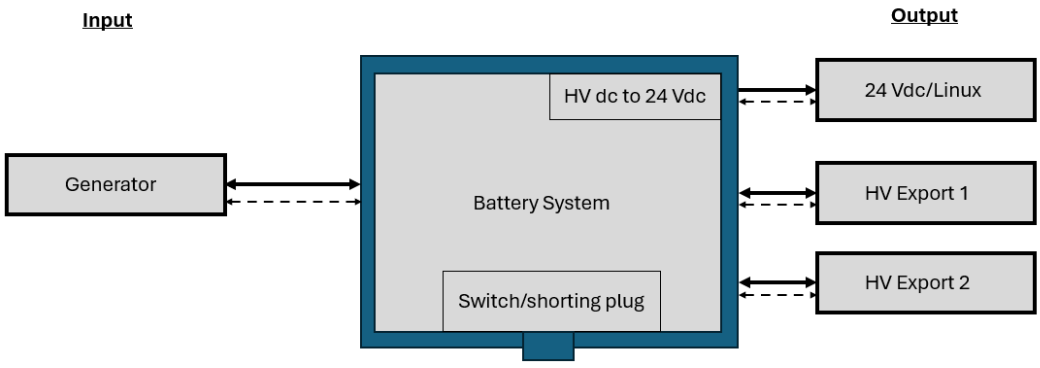

Figure 11: Battery system block diagram. Solid lines represent power transfer and dashed lines represent communications. With the exception of the 24 Vdc export, bidirecitonal power transfer and communications can cross all interfaces. An external shorting plug disables all high voltage contact, isolating the system for safety.

Table 1: Battery system specifications. Note, some values in this table represent design requirements and not the final design values.

Voltage, current, and SOC	
Nominal voltage	270 V
Voltage range	240-307 V
Nominal capacity	9.6 kWhr
Discharge current (continuous)	5 A
Discharge current (peak, < 10 s)	30 A
Surge current (< 0.5 s)	35 A
Charge current	5 A (average), 35 A (peak)
SOC range	80%
Maximum internal resistance	300 mΩ
Chemistry, type, and life cycle	
Chemistry	NMC
Cell type	Vendor-informed form factor
Life cycle	Minimum 1000 cycles
Calendar life	36 months
BMS	
BMS supplier	Vendor recommendation
Communications protocol	CAN
SOC accuracy	<5%
Sleep/low power mode	As needed to minimize power consumption during idle states
Environmental Conditions	
Operating temperature	5–40°C
Temperature control	Passively cooled

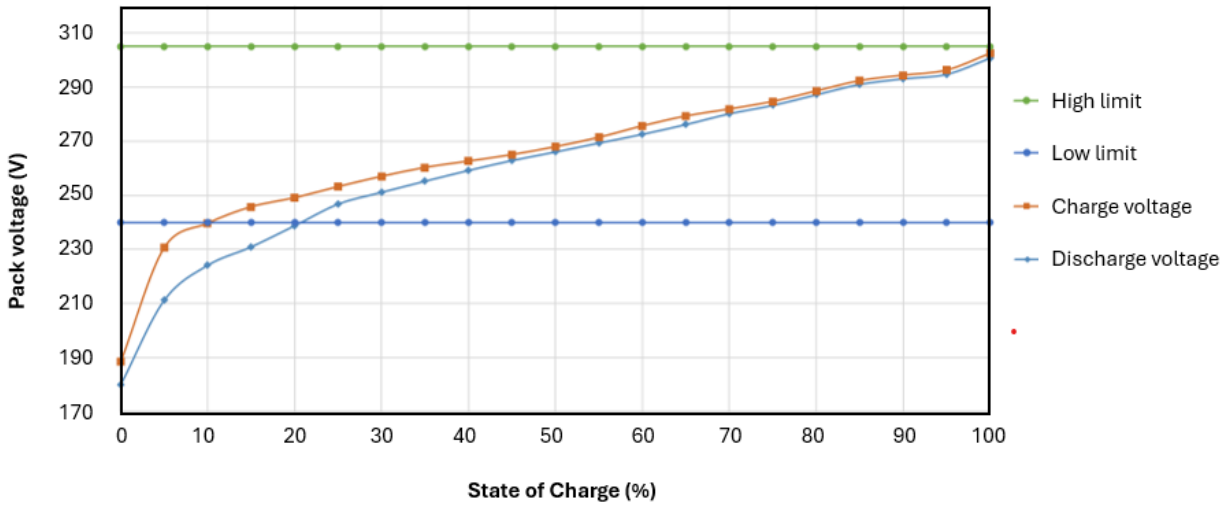

Figure 12: Voltage versus state of charge curve for the NMC battery pack.

If due to high electrical loads or little power production the system requires shut down, the low-voltage disconnect can limit power distribution to the low voltage supply and the system can go into a sleep mode with negligible hotel loads until conditions again allow for power generation. In wave energy applications the system can self-start driven by a wave, but for the tidal energy converter the system can continue to provide occasional power (e.g., on a programmed duty cycle) to the control board to bootstrap the turbine into a state of power production if currents are sufficient. The 1st generator rotor was not capable of self-starting under in situ inflow conditions, which drove this requirement. Testing has since shown that the 2nd generation system can self-start, but this feature could still help to generate power in cases where biofouling of the rotor precludes self-start but power could still be produced.

The cell chemistry is lithium nickel manganese cobalt oxides (NMC). Lithium iron phosphate (LFP) was considered for its stability and ability to accommodate higher thermal loads. However, given the challenging charge/discharge profiles that could emerge in marine energy applications, several vendors recommended NMC because its state of charge slope will allow for easier management of the system's state of charge, which is a significant consideration for the overall safety. Furthermore, given the peak, time-averaged power production of the Turbine Lander and representative profiles of the MBARI WEC, thermal management of the battery was not identified as a significant concern. The recommendation for NMC, as opposed to LFP, was made by several vendors that had reviewer system requirements and charge/discharge profiles for the prototype marine energy converters under consideration. The final pack consists of 576 individual cells distributed among eight parallel packs.

Safety and management features integrated into the pack include a soft start mechanism to mitigate capacitive surge currents at start up, disconnects for all critical power import/export connections, and fault management for standard failure mechanisms. The high voltage power is disabled using an external switch (shorting plug) and the batteries themselves will be shielded with a no touch enclosure. Several operational settings that control when power can be transferred, either from the generator or to external loads based on the stage of charge, will be implemented. A simple block diagram of the system (Figure 11) shows that two-way power and communications will

be enabled for the generator and up to two high voltage (same as nominal battery voltage) power exports. These power exports could be used to parallel the packs to provide additional storage capabilities. In addition, a low voltage (24 Vdc) export will be reserved for providing power to external instrumentation. The battery management systems (BMS) and integrated power supply are powered by the battery itself to avoid the need for external power supplies in the autonomous system. CAN communications protocol is used by all components. While communications between the generator or an external computer are enabled, most of the system management is internal to the battery.

The system itself is chassis-mounted to ease integration into housings for different applications (Figure 13). For the Turbine Lander it is housed in a submersible, atmospheric pressure housing, while in a WEC a submersible housing is not required and a standard enclosure is sufficient. Unhoused, the system's dimensions are approximately 1 m long and sized to fit inside a 12" housing. The battery system is currently in fabrication and will be tested on APL-UW's laboratory dynamometer and simulated in situ deployments using R/V *Russell Davis Light* prior to a future deployment.

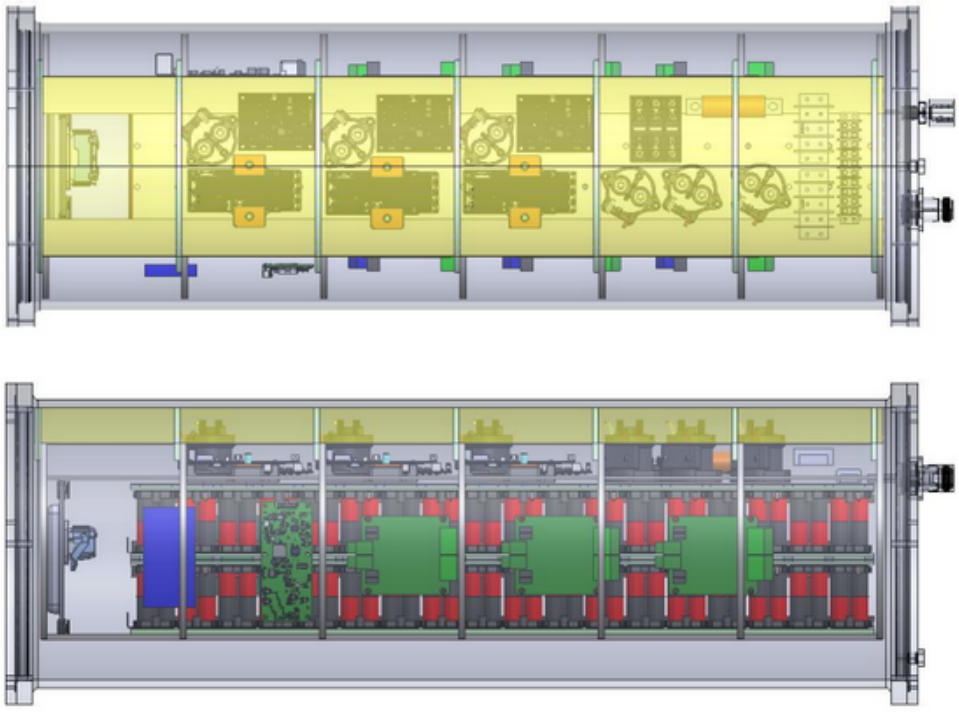

Figure 13: An early mechanical drawing of the battery system, which represents the concept and contains critical components, but varies from the final layout. The housing shown here, which has not been fully designed, is based off a pre-existing housing design.

4 ROTOR AND PTO MODIFICATIONS

This section summarizes design requirements and mechanical modifications made to the rotor design and power take-off in the 2nd generation system. For context, the performance of the 1st generation system, without modifications to the powertrain, is summarized prior to introducing the 2nd generation system.

4.1 First-generation PTO Performance

To compare the 1st and 2nd generation systems, the 1st generation rotor was retested on R/V *Russell Davis Light* in constant speed and constant torque control modes. While these control modes vary in their implementation and limitations, both approaches can yield similar amounts of time-averaged power generation under the same inflow conditions [5]. The only difference between the configurations described in Bassett et al. [2] and the tests reported here is the use of the MPC. The differences in power output and ultimately the water-to-wire efficiency are driven by two factors. First, the results in Bassett et al. [2] used the Siemens equipment and supporting equipment that were powered externally and not represented in the power budget. Furthermore, the reported powers were measured at the motor module and were not representative of the total power budget. In the case of the 2nd generation system, it is designed with integrated power storage in mind so the MPC's output includes the power conditioned to 325 Vdc (nominal), which does not reflect the subsequent decision to lower the operating voltage. In addition, when power is being generated the MPC and microcontroller are self-powered. These conditioning losses and hotel loads were not accounted for in the 1st generation system's performance.

Results for the power delivered to a load (i.e., battery or load dump) under constant speed and torque control are shown in Figures 14 and 15. Based on the discussion in the prior paragraph, the total electrical power output is calculated according to $P_e = P_{load} - P_{source}$, where P_{load} is based on the voltage and electrical current delivered to the resistive load dump and P_{source} is calculated from the voltage and electrical current provided by a dc power supply (0 W when the system is generating enough power to overcome all losses). Thus, this calculation inherently includes all electrical inefficiencies and hotel loads. The water-to-wire efficiency η_{ww} is calculated according to

$$\eta_{ww} = \frac{P_e}{\frac{1}{2}\rho A U_o^3}, \tag{1}$$

where A is the cross-sectional area of the rotor (1.01 m^2), ρ is the water density (1000 kg/m^3), and U_o is the inflow speed.

Calculations of P_e and η_{ww} were made during tests under propulsion at vessel speeds of 2 – 4.5 knots (approximately 1 – 2.25 m/s). Tested conditions were modified at different speeds but the methods incremented rotor speed and control torque in 1 RPM and 1 N-m steps, respectively. Limits and testing at each condition were determined to capture the relevant portions of the performance curves that could be measured under each scheme, but generally avoided conditions at which the system would not be operated due to poor efficiency (i.e., tip-speed ratios well above or below peak values). Each set point was measured for 30 s. Three seconds at the beginning and end of each condition were removed and the remaining 24 s were averaged to calculate relevant performance parameters. Because the vessel's speed is not constant, a Nortek Vector ADV was deployed 1.5 m upstream of the rotor and sampled at 64 Hz. Data were despiked to remove outliers and 5 s moving averages were used to calculate the performance. The reported nominal in the figures is the

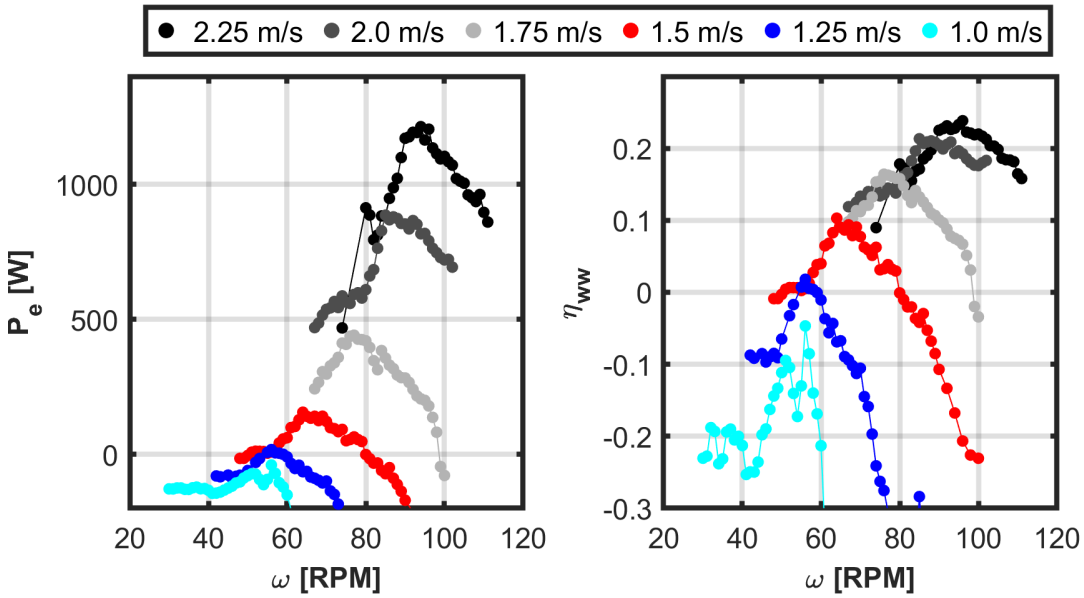

Figure 14: Electrical power generation versus rotor speed at six inflow conditions under constant speed control with the 1st generation rotor.

targeted speed, although in practice individual data points vary by up to ± 0.1 m/s (0.2 knots) of the target. Significant outliers associated with the vessel (rapid turns or speed changes) have been removed. This same acquisition and processing methodology was used to test the 2nd generation rotor and associated modifications.

The performance of the system under speed control (Figure 14) shows that total electrical power generation of the system peaked at approximately 1300 W when rotating at 92 RPM with inflow speeds of approximately 2.25 m/s, while peak power generation was approximately 800 W, 450 W, 200 W, 40 W, and –10 W at inflow speeds of 2 m/s, 1.75 m/s, 1.5 m/s, 1.25 m/s, and 1 m/s, respectively. The decreases in net power production are driven by the available kinetic power density and inherent electrical and mechanical inefficiencies. The water-to-wire efficiencies show that above approximately 2 m/s peak efficiencies exceed $\eta_{ww} = 0.2$. Notably, we estimate that the rotor's coefficient of performance is roughly $C_P = 0.3$, so achieving η_{ww} values approaching 0.25 suggests total mechanical and electrical efficiencies post-power capture by the rotor on the order of 80%. However, the impacts of critical electrical and mechanical losses that don't scale with power capture (e.g., seals and switching losses) degrade system efficiency so overall efficiencies drop to approximately 0 at 1.25 m/s. While little can be done to improve the electrical inefficiencies, Section 4.2 discusses steps taken to decrease the mechanical losses and estimates how these changes would improve the system's performance with the 1st generation rotor.

Optimal implementation of constant speed control schemes required the ability to track inflow conditions, which can vary rapidly over short periods of time due to turbulence. In practice this can be challenging. Alternatively, torque control can provide maximum power point tracking without direct knowledge of inflow conditions[6]. This was not successfully implemented on a regular basis for the first generation Turbine Lander, which is primarily attributed to mechanical inefficiencies in the system that resulted in positive control torques (i.e., power is required to prevent the rotor from stalling) to overcome losses under a broad range of operating conditions encountered in situ. Given

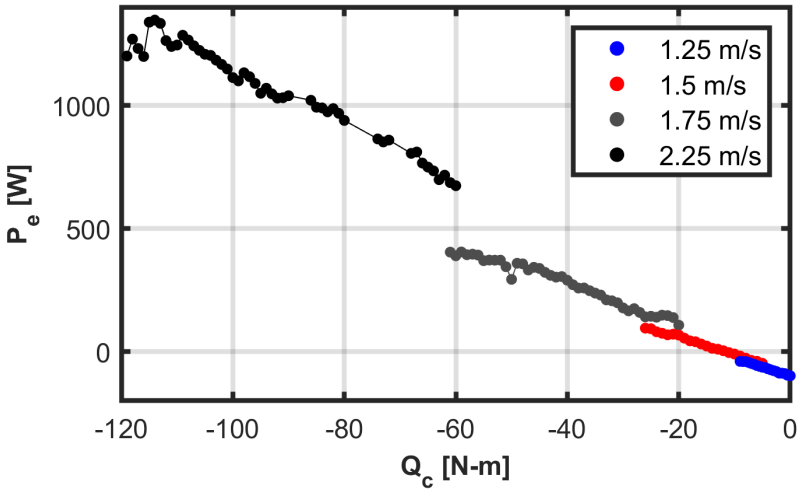

Figure 15: Electrical power versus control torque at four inflow conditions using constant torque control with the 1st generation rotor. Data from a speed of 1.0 m/s is missing because the rotor undergoes required torque reversals at this speed and therefore cannot operate under constant torque control. Data from inflow speeds of 2.0 m/s are missing due to issues with the co-temporal velocity data.

that overcoming these would facilitate other controls schemes, Figure 15 shows the relationship between power generation and constant torque control measured with the 1st generation rotor. Note that the lines in Figure 15 demonstrate a general increase in power production as the control torque increases. The left end of each line corresponds to the highest control torque that can be applied by the rotor without the system stalling. A goal of the 2nd generation rotor design was to shift the Q_c resulting in stall to larger values, thereby creating a larger operating envelope with net positive power production. Identifying modifications to the rotor design to increase hydrodynamic torque generation and minimize the mechanical losses that lead the system to stall, particularly at low speeds, were ultimately significant factors in the system redesign.

4.2 Seal Removal and Bearing Pack Lubrication

During laboratory testing of the 1st generation PTO, the seals were identified to be significant sources of mechanical losses over the operating conditions of the Turbine Lander (approximately 50–120 RPM). Specifically, the PTO included four seals in the design, three of which were determined to contribute significantly to frictional losses in the system. These included the two lip seals to keep the bearing pack from accumulating particulate and saltwater, thereby fouling the bearings. In addition, the factory dust seal on the generator was still installed for the deployment in Sequim Bay. Combined, these seals were estimated to result in frictional losses of roughly 120 to 300 W between 50 and 120 RPM, respectively (Figure 16). During testing to understand the contribution of each seal to the total losses we identified the factory dust seal on the motor to be the most significant contributor to the combined seal losses. This dust seal serves a purpose on a factory floor, but the PTO housing is sealed and has no way of accumulating significant amounts of unwanted particulate. The dust seal was ultimately removed and losses retested (see Bassett et al. [2] for additional details). The removal of the factory seal resulted in a reduction in losses by 50 to 100 W across representative operating conditions, which will increase power generation capabilities considerably.

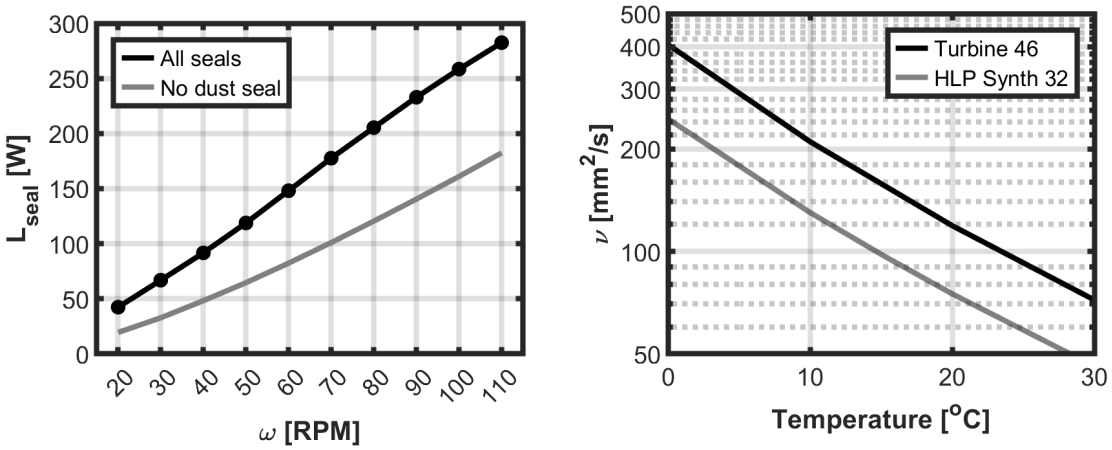

Figure 16: (left) Losses attributed to seals in the PTO. The only difference between the curves is the removal of the motor's factory dust seal. (right) A comparison between the viscosity of the oil used in the 1st generation bearing pack (PANOLIN TURWADA 46, noted as Turbine 46 in the product's name has been changed to Shell PANOLIN S4 Turbine 46) and the 2nd generation bearing pack oil (Shell PANOLIN S4 HLP Synth 32). The values plotted here were approximated from figures included in the manufacturer's product specifications.

The oil in the bearing pack was identified as a less significant source of losses. Without complex numerical simulation that accounts for both the fluid mechanics and heat transfer in the system, predicting these losses is difficult. To simplify the problem we predicted losses by treating the flow in the bearing pack as Couette flow in separate chambers (i.e., proportional of the bearing pack with different geometries). Given that the bearing shaft is stationary and the rotor shaft is spinning. the mean velocity gradient and the total power dissipated by viscous forces can be calculated. We assume any heat generated due to the dissipation is transferred to the environment immediately, resulting in no temperature change, due to the relatively narrow channels occupied by the oil, the highly conductive material of the bearing pack, total mass of the bearing pack, and strong forced conventions by the tidal currents over the system. Under this set of assumptions we estimate viscous losses of 15–40 W across operating conditions using the initial PTO's oil (PANOLIN TURWADA 46). We also noticed in testing that cold conditions resulted in noticeable differences in the resistance of the system being hand-rotated. We thus chose to identify a new, less viscous oil for the bearing pack.

Identifying a new oil for the bearing pack requires a trade-off between several considerations. First, the oil serves as a lubricant for the bearings and prevents the intrusion of seawater and particulate. We are already operating the system with less viscous lubricants than are specified for the PTO's tapered roller bearings, although we have seen no indication that this is leading to premature wear or failure in the bearings. However, this is a factor in favor of more viscous oils. Similarly, a less viscous fluid is more likely to leak from the bearing pack as seals wear and that could lead to particulate intrusion and bearing failure. This also suggests more viscous fluids are beneficial. These factors stand in contrast to the objective of maximizing power generation. To balance these considerations we selected an alternative, biodegradable hydraulic fluid (Shell PANOLIN HLP SYNTH 32), which has a viscosity that is approximately 40% less than the oil used in the initial design (Figure 16). Furthermore, the viscosity of the new oil at 10°C will be similar to that of the prior lubricant at 20°C. While we anticipate that this change in lubricant will

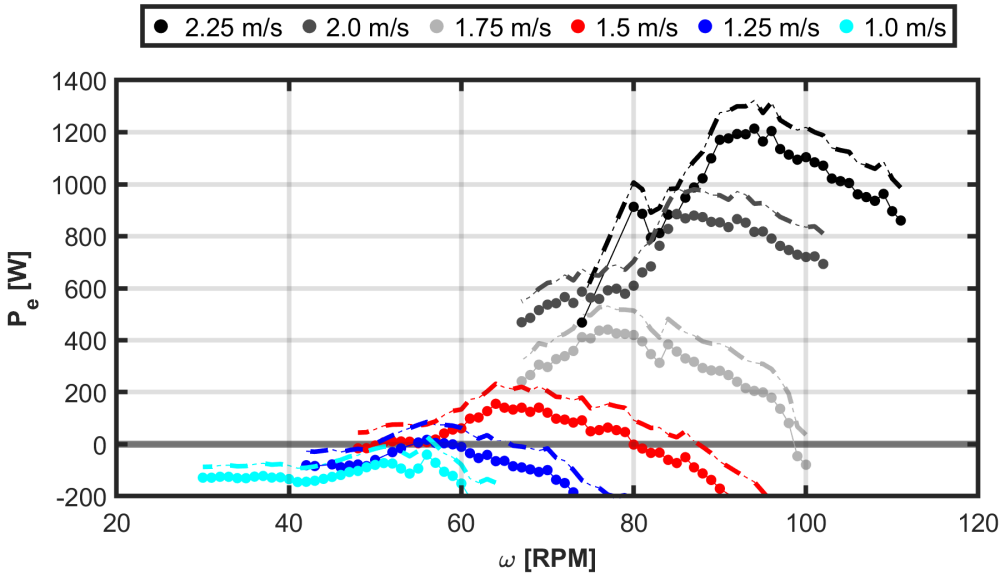

Figure 17: Power versus shaft speed for the 1st generation rotor at the same six inflow conditions shown in Figure 14, but with dust seal removed and viscosity of the bearing pack lubricant reduced. Predictions for additional power production (dashed lines) do not account for changes associated with the rotor redesign.

only yield gains on the order of 5–15 W during operations, we cannot validate this estimate due to our inability to perform a representative test at the target temperatures in the lab.

For reference, Figure 17 includes measured electrical power generation for the first-generation rotor, using the MPC, and adjustments to the power production based on measured and predicted losses from the dust seal and bearing pack oil. The adjustment was made according to $P_e = P_{e,\text{measured}} + 10\omega + 10$ [W], where $P_{e,\text{measured}}$ is the measured power output from the system, 10 ω represents an approximation to the total seal losses based on rotation rate (Figure 16) and the final 10 represents an approximation to the decrease in power dissipated by the oil with the new oil (in practice this is dependent on temperature and rotational speed of the shaft). Across all meaningful operating states we see an increase in net power generation of 60 – 100 W. This represents a relatively minor increase ($< 10\%$) at high speeds, but near 1 m/s this difference is enough to result in net power generation delivered to a battery when the first generation system without these modifications does not produce net positive power. These increases in net power generation could easily result in 10s of additional kW-hr of power generation when added up over a multi-month deployment like that of the Turbine Lander deployment in Sequim Bay.

4.3 Design Requirements and Final Design

Design of a small-scale turbine system is a challenging process with a broad range of under-constrained parameters and unknowns that can only be partially reconciled with archival literature and results from laboratory testing. This was the case with the first-generation turbine and led to the adoption of a design that had been largely tested and evaluated at laboratory scale. While additional information has been published in recent years, many uncertainties remain. The overarching goal of the redesign was to increase power generation as much as can be achieved meaningfully without modifying the Turbine Lander's foundation or bearing pack. An otherwise identical rotor design scaled in size would be a simple way of generating more power, but this approach immedi-

ately violates constraints related to the size and overturning moments that the existing foundation can support and would preclude deployment from APL-UW's most capable research vessel. Thus, we approach the redesign effort from the start with the following constraints:

- The predicted overturning moments from the rotor should not exceed the 1st generation design
- The system's height should not exceed that of the 1st generation system (2.76 m)
- The rotors critical dimensions are flexible, but the final unit should still be small and light enough to be handled by a two-person team
- Redesign efforts should address biofouling concerns identified in the Sequim Bay deployment of the first-generation system while further streamlining components and minimizing exposed hardware that increases drag
- Time-averaged increases in power generation of 50 W were targeted
- The issues identified with the hinge pins in the initial deployment are to be addressed by both strengthening the hinge and providing a more robust method of securing the strut/foil assembly
- Foil preset pitch angles should be modified to optimize power generation
- The four-bladed geometry should not be modified as it represents an appropriate balance between intracycle loads and time-averaged performance.

While these constraints impose size and load constraints on the system, they still leave a large parameter space to consider.

We began the process of rotor design with basic calculations to determine how the rotor's geometry could be modified to increase its overall size without instantaneous forces exceeding the system's overturning moment. Pre-existing calculations associated with the 1st generation system provided the ceiling on the overturning moment. Measurements of system pitch in situ[2] show the system responding to strong inflows and suggest a conservative approach would be to avoid exceeding the overturning moment limits established in the initial design. Thus, because the rotor's height could not be increased, the only option to increase the cross-sectional area was to increase the rotors width and reduce its height. To establish force coefficients used in estimates of the overturning moment we assumed that the four-bladed rotor geometry would be adopted and that the instantaneous force coefficients (C_F) would therefore be consistent with published results for four-bladed rotors with a similar aspect ratio[7]. Extending non-dimensional force coefficients in the literature to predicted overturning moments requires an assumption of peak conditions. Prior Turbine Lander characterization suggested peak performance at tip speed ratios ($\lambda = \omega R/U$, where R is the radius of the rotor) less than 2 while peak, instantaneous inflow conditions consistent with the system's design (i.e., 2.5 m/s) were observed in Sequim. Because force coefficients increase with λ, by assuming higher operational tip speed ratios at peak conditions than we anticipated encountering operationally an inherently conservative assumption is made. We ultimately assumed a peak λ of 2.7, an instantaneous force coefficient of 1.7, and a uniform inflow speed of 2.5 m/s (note that these conditions vary slightly from those in other rotor design analyses, but still represent conservative inputs to the design process). The total overturning moment from the instantaneous force is then

$$F_m = \frac{1}{2}\rho C_F A U_o^2 h_r, \tag{2}$$

where A is the cross-sectional area of the new rotor design, h_r is the mid-plane rotor height above the seabed, and $C_f = 1.7$ is the instantaneous force coefficient, which is dominated by the thrust. The base of the rotor in the design is 1.57 m above the seabed. Iterating on the design rotor

diameter we settled on rotor height of 1.092 m (43") and width of 1.041 m (41") for an aspect ratio of approximately 1.05. Note that laboratory studies suggest that aspect ratios spanning the range of the 1st and 2nd generation rotors (1.4 to 1.06=5) do not have an appreciable impact on performance[8]. Under this set of assumptions the peak moment generated by the rotor would be approximately 13170 N-m (10000 ft-lb). While large, this is consistent with calculations for the 1st generation system.

Hunt et al.[8] suggests that while peak performance of the rotor will not be modified significantly by the change of aspect ratio, the peak tip speed ratio will also be roughly consistent if no other sigificant changes are made (e.g., solidity). This has practical implications for the system's power take-off. Specifically, if the peak tip speed ratio is unmodified, then a larger diameter rotor will spin proportionally slower at peak C_P. To generate equivalent mechanical power this means a proportional increase in torque is generated to account for the drop in rotation rate. Because frictional losses from the seals and bearing pack scale linearly with rotation speed, a decrease in rotational rate will decreases mechanical losses. The scale of the potential decrease in losses associated with a decrease in shaft speed of approximately 20% vary 10 – 40 W across the operating space (Figure 16). In addition, increasing the hydrodynamic torque generation provides an additional buffer relative to the system's baseline mechanical losses, which could decrease rotor cut-in speed and open the door to a broader range of control strategies[2].

With a basic rotor geometry selected, a new design for the foils was needed. We chose a symmetric hydrofoil geometry (NACA 0018) for several reasons. First, more complex foil geometries (cambered blades) impact performance, and this remains an active research area. The NACA 0018 foil is well-studied in the laboratory and was leveraged in the 1st generation turbine design. While overall gains could be made with another geometry, this ultimately represented a safe choice. Two other parameters of interest are the chord to radius ratio and the preset pitch angle. The former relates the size of the foil to the radius of the rotor, while the latter the angle between the centerline of the foil at the quarter chord and the line tangent to the circumference of the rotor. In both cases, the objective was to select a set of parameters that would maximize power generation.

The ideal chord to radius ratio is dependent on many factors, but the primary factor driving our decision was turbine controls in a highly turbulent environment. A lower solidity rotor will operate at a higher tip speed ratio and will have a wider peak in its C_P versus λ curve. This means that turbulent gusts resulting in suboptimal operation will still impact performance, but to a lesser extent. On the other hand, a higher tip speed ratio trades shaft speed for torque generation. Finally, with a predetermined blade count of four, options for modifying solidity are limited to the foil geometry. We ultimately decided to maintain roughly the same chord to radius ratio as the first generation system by using chord lengths of 15.24 cm (6") for a $c/R = 0.29$.

The preset pitch angle of the foils informs overall performance by modifying the foil's angle of attack, which itself changes throughout a rotation. This impacts the balance of lift and drag throughout the rotation cycles creating a trade-off between time-averaged and instantaneous power generation[7;9–11]. On balance, experimental results point to a preset pitch angle of –6° and providing the best balance of power production during upstream sweeps and consumption (drag) on downstream sweeps. This –6° preset pitch angle was the intended target in the 1st generation system although its foils ended up with a preset pitch angle of –9° due to a mistake in the manufacturing process. Experiment results using NACA 0018 foils suggest that the net benefits of the change in preset pitch angle are unlikely to improve C_P by more than a few percent at most, so a preset pitch angle of 4° was selected for the new design. This decision was made in collaboration with other

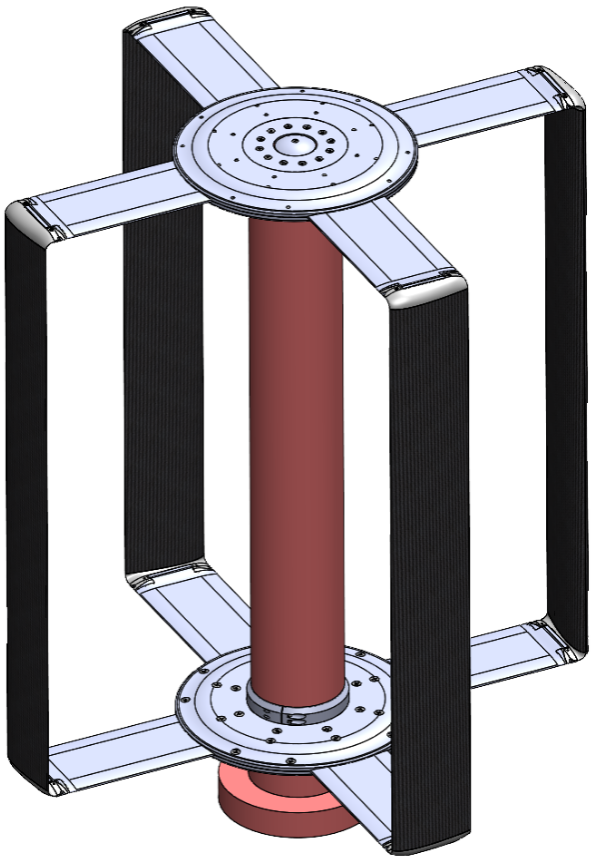

Figure 18: A rendering of the 2nd generation Turbine Lander rotor.

experts at UW based on the results of prior laboratory experiments[7] and discussions related to the objectives of the effort to redesign the rotor.

Having selected rotor and foil geometries, the remaining portion of the design process focused on high-level conceptual tasks (discussed in this section) followed by detailed analysis (Section 4.5) and numerical analyses to ensure that the designs could achieve acceptable safety factors for sustained operations. As with the 1st generation rotor design, the new rotor sought to minimize losses from supporting structures. Experimental work[12;13] shows non-idealized struts or endplate geometries can result in significant parasitic losses that can be detrimental to performance. The 1st generation struts were rounded, but not otherwise streamlined, so the new strut design targeted a modest decrease in their thickness coupled with significant, smoothed surfaces to further reduce their drag coefficients. While no significant reduction in endplate diameters was possible given required coupling with the shaft, the design was streamlined to reduce their blunt faces and recess all of the fasteners and sacrificial anodes to reduce drag and minimize fouling opportunities. The rotor shaft geometry was not modified significantly, although modifications were made to facilitate assembly and to minimize exposed fasteners coupled to the endplates (Figure 18). In addition, the shaft was combined with the standoff to reduce the total part count and then converted into a pressure vessel to minimize biofouling and treepanning of large stock during manufacturing. To facilitate assembly, features were added to the shaft and endplate assembles to ensure rotational alignment between the upper and lower struts and other sub-assemblies (e.g., hinges, pings).

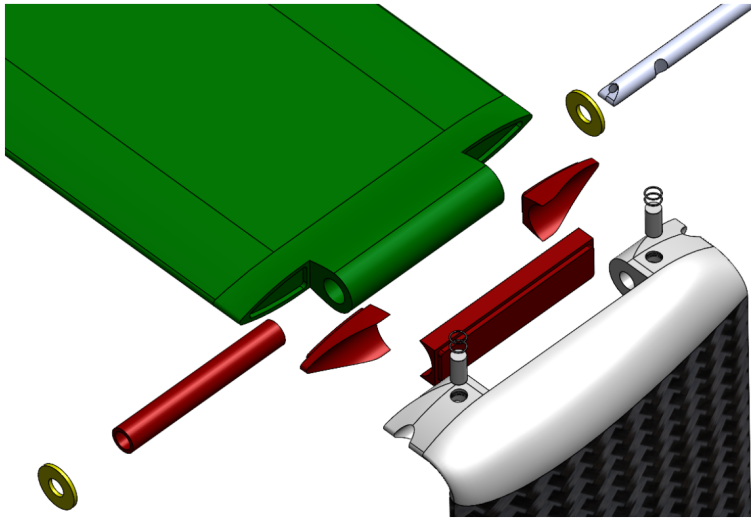

Figure 19: A exploded drawing of the full strut/hinge/foil assembly. The assembly is considerably more complicated than the 1st generation system, but has been redesigned to address several sources of losses and points of failure identified following the deployment of the 1st generation system.

The final significant modifications occurred at the hinges. While there are similarities between the 1st and 2nd generation designs that include the materials and isolation to mitigate mixed metal contact, several modifications were made to address the shortcomings of the first design. In the initial design the pin was held in place by 4-40 titanium set screws. During the deployment two of these screws backed out entirely leading to foil failures while two nearly backed out. To address this the hinge pins will be held in place using retaining pins and retention rings that will be below the shoulder surface to avoid contributions to the drag. Like the strut, the hinges lead and trailing edges have been streamlined to similarly reduce their contribution to drag. During the first deployment eelgrass was caught regularly in the small cracks between the hinges and struts, likely leading to significant decreases in power generation. To minimize this, additional bushings that sit between the hinge and strut will be installed. An added benefit of this component is that post-recovery analysis of the last system suggested that fluctuating forces on the foils had led to some wear of the bushings isolating the titanium pins from the aluminum struts. By closing the gap between the hinges and struts it is possible that motion will be further restricted thereby minimizing wear on the bushings.

4.4 Composite Foil Design

In parallel with working to identify a vendor to fabricate the 2nd generation system's composite rotor foils, a design analysis was performed at APL-UW. The objectives were to finalize the design details and verify that the final product would meet load requirements during in situ deployments.

4.4.1 Design Inputs

To complete the design and analysis, peak operational conditions and physical design constraints were identified. These were driven by prior measurements of inflow speeds and turbulence along with expected rotor performance (C_p versus λ) curves. As previously discussed, the rotor was rescaled and its chord to radius ratio increased modestly. The final foil geometry selected was a

NACA 0018 profile with a chord length of 15.2 cm and a span (length) of 109 cm. A swept radius of 52.5 cm, consistent with the new design, was used. The foil is assumed to be pinned at both ends due to the presence of the hinge joint, with the composite material bonded to the titanium shoulder.

Accounting for the shear in the bottom boundary layer, a peak time-averaged velocity of U_o = 2.4 m/s was selected and, to account for turbulence, a maximum instantaneous velocity of 3 m/s was chosen. Performance characteristics and associated loads are commonly expressed in non-dimensional terms. The tip speed ratio, which relates the tangential velocity of the tips of the rotors to the inflow speed is described by

$$\lambda = \frac{\omega R}{U_o} \tag{3}$$

where R is the rotor swept radius and ω is the angular velocity. Force coefficients, which can be defined to include total forces, thrust or drag forces, and lateral forces, are described by

$$C_m = \frac{F_m}{\frac{1}{2}\rho A U_o^2}, \tag{4}$$

where M is simply a placeholder for an arbitrary force (e.g., $_D$ commonly represents drag). Note, this is a rearrangement of other equations presented in this report (e.g., Eq. 19). At high speeds, the most efficient tip speed ratios are expected to be below 2.0. However, higher λ values as a result of control schemes generated in higher loads and designing to accommodate higher tip speed ratios makes the system more flexible and increases our ability to explore the operational parameter space without concerns for loads. Thus, a maximum $\lambda = 2.5$ was assumed.

It was decided that blade-level force coefficients derived from experimental work with a single-bladed rotor[1] would be used to predict loading, given a lack of other direct measurements with which we could parameterize the analysis. Phase-resolved force coefficients as a function of tip speed ratio are shown in Figure 20. Dr. Abigail Snortland provided additional information from experiment results regarding the location of the center of pressure associated with the highest loads and internal moments experienced by the blades. Our analysis therefore defined the center of pressure to be located 31.5% of the chord length from the leading edge. Dr. Snortland also recommended a scaling factor of 1/0.24 (or 4.17), which is the inverse of the chord to radius ratio, be applied to the single-bladed rotor loading coefficients in Figure 20 to make them applicable to the Turbine Lander rotor.

Lastly, the composite structure was predetermined to be a $\pm 45^\circ$ wrap over a solid unidirectional core. An overview of the general foil geometry in this design is shown in Figure 21.

4.4.2 Load Cases

In speed control, the rotor controller is designed to maintain a selected λ. The theoretically optimal tip speed ratio that maximizes energy production is hereby referred to as λ^*, and for the 1st generation Turbine Lander was approximately $\lambda = 2$, which we expected to be similar for the new rotor. Under speed control, the inflow conditions are measured in situ and are passed to the generator controller, which can modify the resistive torque applied to the rotor to maintain a chosen λ, ideally λ^*. While alternative controls can be implemented, the assumption of speed control provides adequate constraints for the design. In practice, several issues cause the actual conditions and performance to differ in situ from controlled studies:

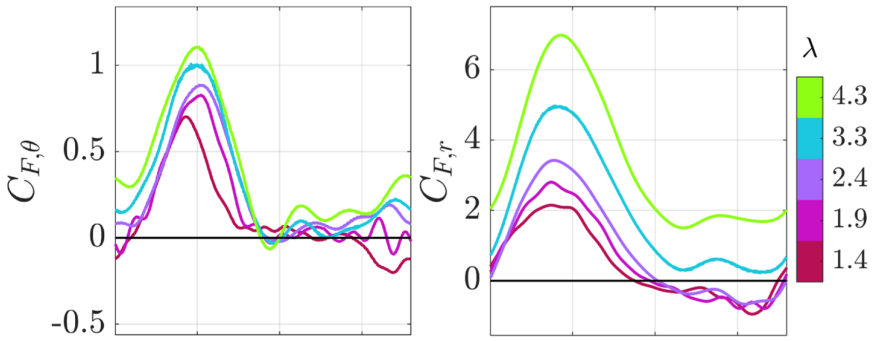

Figure 20: Phase resolved blade-level tangential ($C_{F,\theta}$) and radial ($C_{F,r}$) load coefficients as a function of λ for a single blade rotor from experimental data in Snortland et al.[1] The values are before scaling required to make them applicable to the Turbine Lander rotor.

- Real-world conditions vary continuously in time and this variability is observed in both direction and speed. Thus, the instantaneous velocity experienced by the rotor at any given time can vary.
- Measurements of inflow cannot provide perfect controls inputs. This is attributed to many factors including the volume sampled by the measurement device, impacts of turbine superstructure, and turbulent structures not readily captured by the sampling.
- Even with ideal inflow conditions, controller implementation will be imperfect and will not maintain the ideal rotational speed without variability.
- Significant vertical shear may be present (e.g., Bassett et al.[2]) that would result in variations in tip speed ratio along the span of the rotor, meaning that the optimal tip speed ratio is only occurring over a portion of the rotor.

The resulting deviations in inflow speed and operating state are accounted for in two load cases.

Load Case 1 - Expected Fatigue Conditions: To evaluate fatigue, assume $U_o = 2.4$ m/s and $\lambda = 2.5$. Following Eq. 3, this yields a rotation rate of 109 RPM. Load case 1 (LC1) should be used for considerations of fatigue if, as is the case for this project, the rotor should survive for the design life of the system without requiring velocity data other than the maximum U_o. Of course, in practice, currents seldom reach this design condition and this choice therefore represents a conservative set of assumptions. However, by doing so, this allows for additional flexibility to consider other energetic

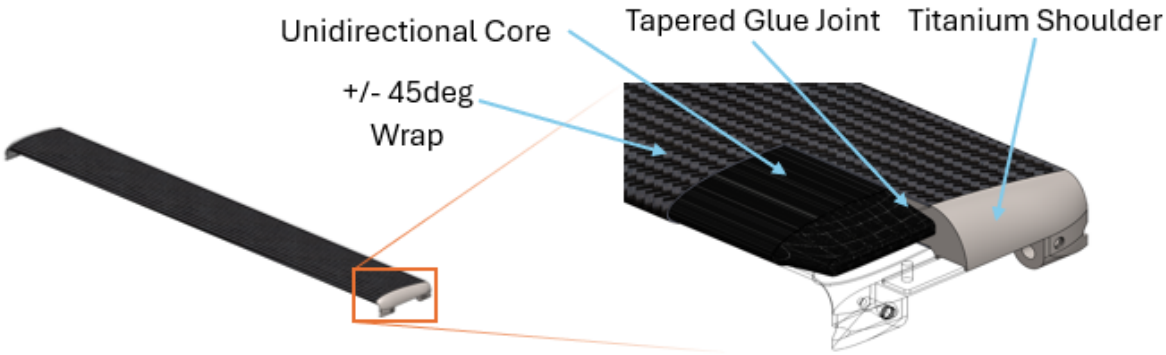

Figure 21: The anatomy of the composite foils for the 2nd generation rotor.

sites for a potential deployment. Furthermore, having selected a conservative tip speed ratio for design calculations, $\lambda = 2.5$ instead of $\lambda \simeq 2.0$, the forces are assumed to be higher ($C_{F,r} \simeq 3.8$ versus $C_{F,r} \simeq 3.0$), where the subscripts denote the radial forces.

Load Case 2 - Maximum Instantaneous Velocity: Here, assume $U_o = 3$ m/s and $\lambda = 2.0$ (roughly the targeted tip speed ratio). This yields a rotation rate the same as Load Case 1, 109 RPM. Load case 2 (LC2) represents a turbulent gust occurring while the rotor is spinning at the design speed for LC1. It assumes that the controller does not change the rotor speed to account for the gust and therefore the λ drops in comparison to LC1. The drop in TSR and corresponding drop in $C_{F,r}$ is more than offset by the increase in water velocity due to the force having a U^2 dependence, resulting in LC2 representing the peak instantaneous loading on the blades.

Radial loads on the blades are calculated following

$$F_r = \frac{1}{2}\rho C_{F,r} c h U_o^2, \tag{5}$$

where c is the chord length and h is the blade span. Loads on the blade tangent to the rotor are calculated similarly:

$$F_\theta = \frac{1}{2}\rho C_{F,\theta} c h U_o^2. \tag{6}$$

Table 2 summarizes the design and expected load cases and associated system states. The design load cases have approximately 30% greater radial loads and 10% greater tangential loads than the expected operational load states.

Note that these values must still be translated into the blade structure reference frame, as the preset pitch angle of 4^o also creates a corresponding angular offset between the radial and tangential loads in Table 2, and those in the chordwise and transverse-to-chord directions that form the natural coordinate system for the blade. Defining the preset pitch angle as α_p, the chordwise and transverse blade loads are calculated as:

$$F_{chord} = F_\theta \cos(\alpha_p) + F_r \sin(\alpha_p), \tag{7}$$

and

$$F_{trans} = F_r \cos(\alpha_p) - F_\theta \sin(\alpha_p). \tag{8}$$

Applying these to the values in Table 2 results in the following design and expected loads in the blade coordinate system (Table 3).

4.4.3 Material Selection and Properties

Material options for the core and wrap were limited to the selected composite manufacturer's (ALLRed and Associates, Elbridge, NY) experience and stock-on-hand to avoid procurement delays and minimize risks caused by using unfamiliar materials. ALLRed maintains stock of Mitsubishi Composites/Newport High Performance Materials products, namely prepregs utilizing the NHM PL301 resin system. This led to basing the final calculations on manufacturer TDS values for the following products:

- Core: Standard Modulus Uni-directional Prepreg, 35% RC using Grafil 34-700 fiber
- Wrap: 3k 2x2 Standard Modulus Twill Prepreg, 42% RC using Pyrofil TR30S fiber
- Bonding Adhesive: Pro-Set ADV-176 / ADV-276

Table 2: Design load cases and corresponding expected operational load cases. The expected load cases represent how a system may be intended to operate while the design load cases represent conservative assumptions representing realistic scenario possibilities.

Case	U_o [m/s]	λ	Rotor speed [RPM]	Scaled $C_{F,r}$	F_r [kN]	Scaled $C_{F,\theta}$	F_θ [kN]
Design LC1	2.4	2.5	109	14.9	7.4	3.7	1.9
Design LC2	3.0	2.0	109	12.1	9.4	3.5	2.7
Expected LC1	2.4	1.9	83	11.5	5.7	3.4	1.7
Expected LC2	3.0	1.5	83	9.6	7.5	3.0	2.3

Table 3: Design load cases and corresponding expected operational load cases in the blade structure coordinate system.

Case	F_{chord} [kN]	F_{trans} [kN]
Design LC1	2.4	7.3
Design LC2	3.4	9.2
Expected LC1	2.1	5.6
Expected LC2	2.8	7.3

Material Properties - Arriving at Design Values: Analysis of composite structures is best performed using material property data derived by testing samples of composite material of the relevant fiber and resin. The samples should be produced using the same manufacturing controls and methods that will be utilized in making final parts, and the samples should be conditioned via aging, heating/cooling, and exposure to the surrounding medium of interest, which in this case is saltwater.

Such an extensive testing and validation program is not feasible for many projects, including the current effort. In lieu of such a program, mechanical properties as reported in product technical datasheets are used as the foundation of this analysis. The values from these data sheets imperfectly match the expected conditions of interest, and do not account for voids or variations in fiber volume fraction. For example, the TDS for the selected wrap material reports a limited number of mechanical properties at only two conditions: (1) At room temperature and ambient humidity and (2) after 14 days of soaking in 71°C water of unspecified salinity. The condition of interest is expected to be 10°C and several months of soaking in seawater. Mechanical properties from the datasheets must therefore be adjusted to conservative values as part of the analysis process.

Composite materials have significantly more complex behavior than isotropic linear-elastic metals operating within their linear regime. Classical Laminate Theory (CLT) provides a framework to determine structure-level behavior of composites. The basic workflow includes:

1. Measure the mechanical properties of the constituent materials (e.g., fibers and resin).

Units: [Pa]	Tensile Modulus	Tensile Strength Basis	Tensile Stress LC1 Allowable	Tensile Stress LC2 Allowable	Shear Stress LC1 Allowable	Shear Stress LC2 Allowable
Core	1.28E+11	1.37E+09	9.10E+08	1.23E+09	9.72E+07	6.48E+07
Wrap	6.41E+10	6.14E+08	4.09E+08	5.52E+08	2.83E+07	1.88E+07
Adhesive	1.93E+09	3.67E+07	2.45E+07	3.31E+07	1.99E+07	1.32E+07

Figure 22: Material properties for the foil materials, including the uniaxial core, wrap, and adhesive.

2. Calculate the mechanical properties of a single ply (lamina) of the composite material, with the values based on the coordinate system derived from the fiber direction.
3. Calculate the lamina mechanical properties based on the coordinate system used in analysis of the structure.
4. Determine the stiffness behavior of the structure in the structure's coordinate system.
5. Calculate deflections in the structure as a result of the loading conditions.
6. Calculate the resulting strains in the structure at the points of interest for each ply.
7. Calculate the design margins either in terms of strain or stress allowables.

The technical datasheets for the composite materials used in this project report strength and elastic properties from testing of unidirectional and ±45° laminates. Material properties used for design purposes are shown in Figure 22. For bending stresses, the compressive strength is used as the design basis because the selected materials have lower compressive strength than tensile strength, which is typical for polymer matrix composites. Allowable stresses are 60% of the design basis strength value for LC1, and 90% for LC2. The design margin is defined by

$$\text{Stress Design Margin} = \frac{\sigma_{allow}}{\sigma_{LC}} - 1, \tag{9}$$

where σ_{allow} is the allowable design stress driven by material properties (Figure 22) and σ_{LC} is the maximum resultant stress evaluated for a given load case. Stress design margins greater than 0% indicate that the maximum stress occurring in the part during the design load case will not exceed the allowable stress. Stress design margins less than or equal to 0% indicate that the actual stress is expected to be greater than allowable stress in the design. Stress design margins greater than 0% indicate that the maximum stress occurring in the part during the design load case will not exceed the allowable stress. Stress design margins less than or equal to 0% indicate that the actual stress is expected to be greater than allowable stress in the design.

4.4.4 Structural Analysis and Stress States

A NACA 0018 foil has no camber and a foil thickness that is 18% the thickness of the chord length. Thus, for the 15 cm long foil maximum foil thickness is 2.7 cm. In our design the wrap occupies a 1 mm thick shell inside the defined blade profile, equating to four plies at an estimated fiber volume fraction of 60% (the cured ply thickness is 0.25 mm per material supplier guidance). The solid core material occupies the entire volume inside the wrap. While structurally unnecessary, and suboptimal from a weight perspective, a solid core simplifies the fabrication process, enhances the rigidity of the foils, and is a better solution for a small-batch production run. We estimate the mass could be reduced by 20% while still meeting design requirements, if necessary, but note that the heavier foil and the associated larger rotational moment of inertia may have benefits to

the system that include decrease velocity fluctuations associated with turbulence and intracycle hydrodynamic torque generation.

For analysis, we assume the center of the pressure is located 31.5% of the chord distance from the leading edge (Figure 23). This value represents the midpoint between the chord positions of the peak and average and peak center of pressures (personal communication with A. Snortland). Similarly, the location of the foil's centroid and maximum thickness are labeled in Figure 23.

Structural analysis focused on five features of the foil design: Bending stress in the core, transverse shear in the core, torsional shear stress in the exterior wrap, shear stress in the titanium shoulder joint adhesive, and shear stress in the bond line between the leading and trailing edges of the core. Several key assumptions and simplifications were made in support of the analysis. These include:

- The foils are mounted in a hinge that prevents bending moments at the hinge joint. This, which is a protection measure for the glued joint connecting the titanium shoulder to the carbon fiber, shifts the point of maximum bending stress to the center of the foil.
- Bending stresses about the thickness axis are negligible.
- Torsional moments are supported only by the $\pm 45°$ wrap while bending and transverse shear loads are supported by the unidirectional core. This assumption is based on the relative stiffnesses of the core and the wrap. In theory, the core can also support the anticipated torsional stresses by itself but prior experience and recommendations from industry experts lead us to include the wrap.
- The joining of leading and trailing edge sections of the core, while actually occurring over a keyed profile to increase the contact area (Sec. 4.6.1), is analyzed as occurring over a flat joint.

Stress State 1 - Core in Bending:
As stated previously, the core experiences its maximum bending stress at the center of the span. By making the conservative assumption that the core is solely responsible for the supporting bending loads, a simple beam bending analysis can be performed. Because the core is made up of fibers aligned in the same coordinate system as the structure, lamina level strength properties are used

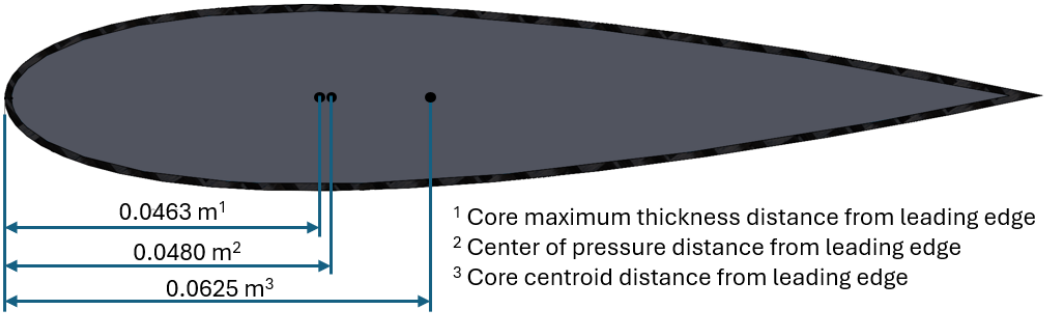

Figure 23: Foil profile showing the chordwise locations of the maximum core thickness, exterior aerodynamic center of pressure, and the centroid of the core.

Figure 24: The blade structure coordinate system drawn on foil rendering.

directly to determine the bending stress design margin of the core according to

$$\sigma_{max,corebending} = \frac{M_{max}}{S_{core}}, \tag{10}$$

where S_{core} is the section modulus as determined from the design drawings and M_{max} is the maximum bending stress at the center of the beam for a uniformly loaded, simply supported beam[14]:

$$M_{max} = \frac{F_{total}S}{8}, \tag{11}$$

where F_{total} represents the total radial force on the beam and S represents the blade span. Table 4 shows the maximum predicted stresses and the design margins for the core in bending, which exceed 400% for both load cases.

Table 4: Stress estimates for stress state 1 (core in bending)

[Pa]	Maximum Stress	Design Margin
Load Case 1	1.44E+08	468%
Load Case 2	1.89E+08	576%

Stress State 2- Transverse Shear in Core:
Transverse shear in the core is calculated at the hinge joints along the chord of profile. This is conservative as the composite tapers into the u-joint before reaching the hinge location. As with the bending stresses in the core, the unidirectional core is assumed to be solely responsible for resisting the transverse shear stress, and the lamina shear stress allowable can be used directly due to the alignment of the fibers with the span of the structure. The maximum predicted shear stress is calculated according to

$$\tau_{max,coretransverse} = \frac{F_{transverse}/2Q}{I_{x,core}c_{core}}, \tag{12}$$

where Q is the first area moment of inertia of one half of the blade cross section about the chord line, $I_{x,core}$ is the second area moment of inertia for the core cross-section, and c_{core} is the chord length of the core (as opposed to the chord length of the full blade). Table 5 shows the predictions

for maximum transverse shear stresses in the core and the design margins, which exceed 400% for load cases 1 and 2.

Table 5: Stress estimates for stress state 2 (transverse shear in core)

[Pa]	Maximum Stress	Design Margin
Load Case 1	1.01E+07	477%
Load Case 2	1.27E+07	588%

Stress State 3 - Torsional Stress in the External Wrap:
The wrap is analyzed as supporting the torsional load by itself, although analysis suggests that the core will support a noticeable fraction of the torsional load with significant margin. Torsion is created by the separation of the geometric center of the blade and the aerodynamic center. The hinge joints act to resist the induced torsion. The maximum torsion occurs at the hinge joints, and is calculated according to

$$T_{max,wrap} = \frac{F_{total} l_{cop}}{2}, \tag{13}$$

where l_{cop} represents the distance from the center of pressure to the centroid of the core (see Figure 23). This torsion is then used to calculate that maximum torsional stress in a thin, closed wall shell following Young et al. [15], Table 10.1, cases 19–21. Thus, the maximum torsional shear stress is described by

$$\tau_{max,wrap} = \frac{T_{max,wrap} C_{wrap}}{K_{wrap}}, \tag{14}$$

where C_{wrap} and K_{wrap} are defined as

$$\begin{aligned} C_{wrap} &= \frac{D}{1 + \frac{\pi^2}{16A^2}} \left[1 + 0.15 \frac{\pi^2 D^4}{16A^2} - \frac{D}{2r} \right] \text{ and} \\ K_{wrap} &= \frac{4 I_x}{1 + \frac{16 I_x}{A c_{core}^2}}. \end{aligned} \tag{15}$$

The previously undefined terms in C_{wrap} and K_{wrap} are related such that the largest inscribed circle associated with the wrap has a diameter of D, a cross-sectional area of A, and a radius of curvature r at the location where the inscribed circle "contacts" the wrap.

The predicted stresses and design margins for torsional stresses in the external wrap are shown in Table 6. In contrast to the core, which has large design margins, the wrap itself appears more marginal. We note, however, that an underlying assumption is these calculations was that the wrap itself supports the full torsional load. In practice, we expect the core to support a considerable portion of the loads thereby providing adequate designs margins.

Table 6: Torsional stress estimates and design margins for stress state 3 (torsional stresses in the external wrap).

[Pa]	Maximum Stress	Design Margin
Load Case 1	7.90E+06	115%
Load Case 2	9.97E+06	155%

Stress State 4 - Shear Stress at the Titanium Shoulder Glue Joint:
The shoulder glue joint (Figure 25) must be able to support transferring the internal moments

of bending and torsion. The joint is tapered to avoid stress concentrations due to the change in stiffness between the titanium and composite materials.

This tapering makes precise hand calculations challenging, and the lack of a stress-strain behavior model for the adhesive complicate FEA. We adopt a simplified model of the joint where the internal moments are assumed to be transferred between the composite core and titanium shoulder by two lap joints conservatively located within the true profile of the lap joint. In this case, the lap joint planes are assumed to be separated by 7.4 mm (or 3.7 mm above and below the chord line, are 92.7 mm long (parallel to the chord), and 25.7 mm wide (parallel to the span). The joint is assumed to transfer internal moments purely by shear loads on the face of these lap planes. See Figure 26 for a visualization of simplified adhesive joint planes and load directions. Note that the moment being transferred in the joint is assumed to be the one where the core begins its taper.

The shear force as a result of internal moments (either from torsion or bending) is calculated by

$$F_{shear} = \frac{M}{l_{space}}, \tag{16}$$

where l_{space} is the gap between the lapjoint planes. The maximum in-plane shear force is the vector sum of the torsion and bending moment shear forces (M), which, when divided by the lap joint plane area, results in the maximum shear stresses and design margins based on the adhesive stress allowables (Table 7).

Table 7: Stress estimates for stress state 4 (titanium shoulder glue joint)

[Pa]	Maximum Stress	Design Margin
Load Case 1	9.60E+06	24%
Load Case 2	1.21E+07	48%

Stress State 5 - Shear Stress at the Core Bond Line:
The unidirectional core was constructed by bonding two unidirectional laminates together. The bond line was a keyed profile transverse to the chord at the thickest point of the core cross section

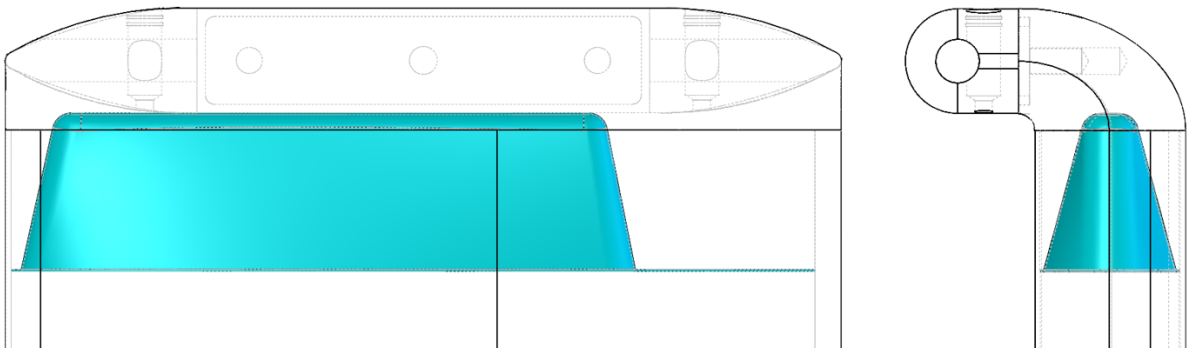

Figure 25: Detailed views of a foil's tapered adhesive shoulder joint with solid color showing the joint extending from the machined core of the composite blade.

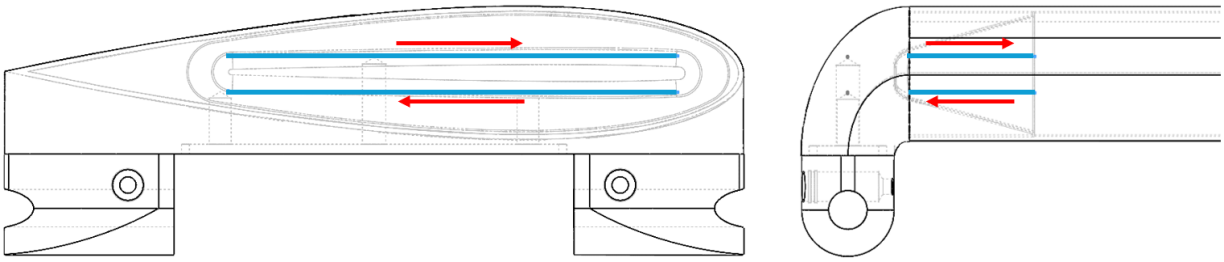

Figure 26: Detailed views showing position of simplified lap joint planes (blue) and shear forces (red) due to torsion (left) and internal bending loads (right).

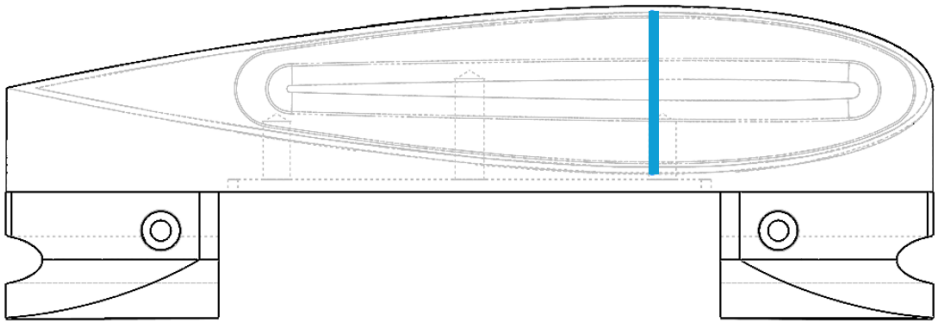

Figure 27: Detailed views of showing position of simplified core bond line (blue).

(Figure 30). The center of pressure nearly perfectly overlaps the bond line (Figure 23), which minimizes the internal moment about the span experienced by the bond line. Thus the bond line primarily needs to withstand internal shear loads resulting from the chordwise blade loads. As the keyed bond line profile increases, so does the surface area through which the internal shear loads are transferred, ensuring that cured laminate participates in the load path. As stated previously, the joint is analyzed as though the the bond line is flat through the thickness of the core profile.

Maximum shear stress in the bondline is calculated similarly to that of the core resulting from transverse loading:

$$\tau_{max,bondline} = \frac{\frac{F_{chordwise}}{2} Q_{halve}}{I_{y,core} t_{core}}, \tag{17}$$

where Q_{halve} is the first section moment of the area on one side of the bond line, $I_{y,core}$ is the second area moment of inertia about the bond line axis, and t_{core} is the thickness of the core at the bond line. Utlilizing strength allowables for the adheisve, the resulting stresses and design margins for the core bond are shown in Table 8.

Table 8: Stress estimates for stress state 5 (core adhesive bond line)

[Pa]	Maximum Stress	Design Margin
Load Case 1	1.81E+06	559%
Load Case 2	2.30E+06	678%

The overall limiting feature in the foil design is the titanium shoulder glue joint (stress state 4) in the fatigue loading condition (LC1), with a design margin of 24%. This design margin should still be conservative based on the geometric simplifications involved in its calculation, but it indicates

that it is likely the first point of structural failure in the design. Due to the solid, undirectional core, the foil has considerable design margins as was expected when, for manufacturing purposes, we opted not to use a design with a foam core. These outcomes were expected throughout the design process because minimizing hydrodynamic losses by maintaining a small, streamlined joint at the interface between the foils and struts was a design priority.

4.5 Other Supporting Analytical and Numerical Analysis

4.5.1 Rotor, Foil, and Pin Force Calculations

For mechanical design of the blades we needed to predict the individual forces generated by the foils during operations. We adopted the framework used by Snortland et al. [11], which infers blade level forces from laboratory tests using load cells and rotors with a single foil. Furthermore, by operating the rotor without the foil, contributions to measured forces from other parts of the rotor can be accounted for in force estimates. This approach was used given the lack of blade-level forces measurements available in literature. The non-dimensionalized force coefficient, which is a function of the tip speed ratio and angular position of the foil (θ) is described by

$$C_F(\lambda, \theta) = \frac{F^*(\lambda, \theta)}{\frac{1}{2}\rho U_o^2 D S}, \tag{18}$$

where U_o represents the undistributed inflow speed upstream of the rotor, D is the diameter of the rotor, and S is the blade span (height) of the rotor. Here, F^* represents the radial forcing from the fluid on the foil. Because C_F is generally scaled by the area of the rotor, this force measurement must be rescaled such that D is replaced with c, the chord length. The geometry testing in Snortland et al. [1] had a chord-to-radius (c/R) ratio of 0.24, which is slightly smaller than the 2nd generation rotor's design in which $c/R = 0.29$. Lacking good alternatives, we reference the intracycle loads reported by Snortland et al. to predict the rotor's performance under field conditions.

Input values for C_F required an assumption about operational tip speed ratios and inflow conditions. To represent a worst-case scenario we assume a uniform inflow speed of $U_o = 2.5$ m/s and $\lambda = 3.3$. Under these conditions the peak predicted intracycle force coefficient is roughly $C_F = 5.5$. This is much higher than $C_F = 3.5$ (λ =2.4) or $C_F = 3.0$ (λ =1.9), which would be expected closer to the rotor's targeted operating state for optimum efficiency, but was selected to represent a conservative value. Rearranging Eq. 18 to solve for F^* yields forces of approximately 10 kN for $C_F = 5.5$ and 6.5 kN for $C_F = 3.5$. Although we also calculated loads for velocity profiles by performing piecewise integration with loads in 0.1-m sections, the assumption of uniform flow at the design conditions represents a conservative design choice, particularly given that the rotor is subjected to varying velocity profiles due to turbines. These peak conditions are used in analytical and numerical studies of other system components including the foils, struts, hinges, and pins.

4.5.2 Bearing Loads and Shaft Displacement

The shaft size of the 1st generation Turbine Lander was based on a decision to restrict the reflection of the shaft to less than ≈ 3 mm under peak loads and to allow the system to be lifted by the shaft during deployment. This deflection limit was chosen primarily to decrease the likelihood of shaft displacement causing the seals to leak oil from the bearing pack. Given that no significant

leakage was identified following the 1st generation system deployment, we chose to maintain the same shaft diameter. To identify the expected peaks in intracycle force coefficients for a four-bladed rotor we referenced Hunt et al. (2024)[7]. Taking a conservative approach we selected a peak force coefficient of $C_F = 1.7$ assuming λ =3.5. Note this is much lower than the single blade values presented above for a single foil and reflects the decision to calculate bearing loads and shaft displacements based on existing laboratory measurements for the forces generated by four-bladed rotors. For the new design, the peak force on the rotor is then

$$F = C_F \frac{1}{2} \rho U_o^2 A, \tag{19}$$

where $A = 1.14$ m^2 and $\rho = 1030$ kg/m^3. The predicted peak for the 1st generation system, which did not overturn during its deployment, was 7.9 kN, a value that accommodates $U_o < 2.8$ m/s despite some conservative assumptions. This value exceeds the peak instantaneous currents measured in situ so the increase in geometry was determined to not violate restrictions on the system's overturning moment.

To avoid redesigning the bearing pack the new rotor geometry had to be accommodating. The initial design consists of two tapered-rolling bearings separated by 14 cm. The shaft is cantilevered above the bearing pack and, for these calculations, it is assumed that total load is well-represented by a single force applied at the mid-span of the rotor approximately 1.2 m from our origin (defined as the location of the lower bearing; Figure 28). The tapered roller bearings were Timkin 52400/52618, which can accommodate radial forces of 59.7 N (13000 lbf) and 223 kN (50200 lbf) for 90 M and 1 M revolutions, respectively. The ratings for dynamic thrust at 90 M revolutions are 47.1 kN (10600 lbf). To estimate bearing loads we calculated the distribution of loads using the peak instantaneous forces from the rotor and the driveline geometry (Figure 28).

Using the 7900 N force generated by the rotor with additional loads of approximately 1300 N generated by the shaft and its standoff (a conservation number due to the velocity profile and velocity deficit within the rotor due to the foils), the total maximum thrust load is about 9200 N. Using the free body diagram, the peak thrust loads are well below what can be accommodated by the bearings while maximum estimated peak radial loads are 33.9 kN (7621 lbf) and 45.2 kN (9464 lbf), respectively. This falls well within the design limitations of the bearings for high cycle counts, especially given that currents at sites suitable for deployment are expected to be well below the design conditions the majority of the time.

Numerical analysis was also performed on critical components of the assembly using ANSYS. Struts, hinge pings, shaft deflections, and endplate/fastener assemblies were modeled using peak predicted loads consistent with those described above to verify component capabilities. Additional details regarding the implementation of the finite element models are not presented here, but Figure 29 includes examples of modeled stresses on components including the end plates, struts, fasteners, hinges, and hinge pings under predicted peak loads. Critically, across the different loading scenarios the minimum safety factors identified in the numerical analysis were on the order of 1.25 for the fasteners and 2 (or greater) for the remaining components.

4.6 Rotor Material Selection

The materials selected for the new rotor design were driven by several factors. First, when possible, we sought to use material consistent with the 1st generation design. Adjustments were

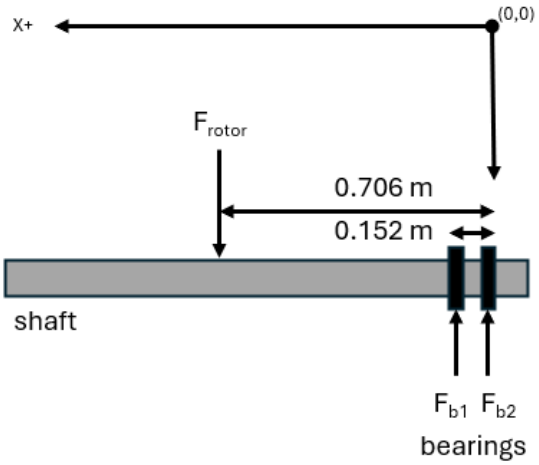

Figure 28: The free body diagram used to estimate the bearing loads. Note that the distance between the rotor force (F_{rotor}) and the bearings reflects the length of the rotor shaft and the shaft's standoff.

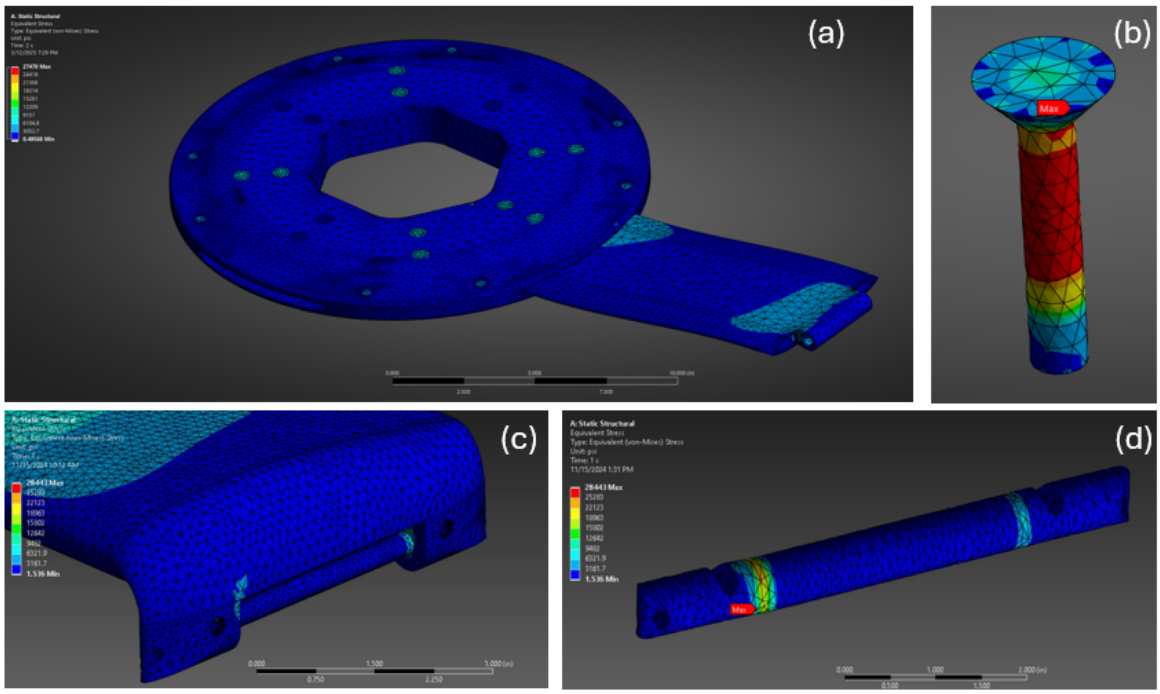

Figure 29: Sample outputs of finite element analysis performed on the endplates, struts, hinges, and fasteners. These analyses supported the assessment that the designs should be able to accommodate loads during operation and highlight the endplate fasteners and hinge pins (particularly at the point of contact with the hinges) as the components subject to the highest stresses relative to their strength.

made only when necessary to accommodate changes to the design. All materials selections were driven by the analysis in Section 4.5 or by general considerations related to corrosion. The rotor

shaft, endplates, and struts are made from hard-anodized 6061-T6 aluminum, which were machined at APL-UW. While the use of 6061 fasteners would have been preferred, analysis suggested that they were not sufficient for the application. This left several options:

- Use larger, exposed fasteners that could not be easily recessed
- Reconsider the design that streamlines the endplates
- Identify an alternative metal with sufficient strength while minimizing corrosion risk

We ultimately decided to use 7075 aluminum fasteners in the end plates. Due to its proximity to 6061 in the galvanic series the risk of corrosion is minimized, although efforts will be taken to isolate the exposed surfaces of the fasteners during saltwater deployments. The hinges, pins, and associated hardware were T5 titanium and procured from a outside vendor as were the bushing (acetal) and washers (PEEK). The foils, which are described in detail in Section 4.6.1 were solid core composite foils with a 1 mm wrap. The remaining materials were unchanged and are described in the summary of the 1st generation system[2].

4.6.1 Foil Fabrication and Testing

The custom composite foils for the 2nd generation unit has solid cores formed from Mitsubishi Grafil 34-700 fiber with 35% resin. This standard modulus, uni-directional, prepreg composite was determined to have sufficient strength to support a foam core or solid core foil; we elected to use solid core foils to simplify manufacturing. The solid core is wrapped with 1 mm of Mitsubishi TR30S 3k standard moduls twill preprep with 42% resin content. Newport Materials 301 resin was used for the wrap and core. Pro-set ADV-176 and ADV-276 adhesives were used as the bonding agents for the individual components.

The foils were manufactured by Allred & Associates (Elbridge, NY), who worked with APL-UW to develop a manufacturing plan that could meet the requirements of the project. Several pictures taken from different portions of the manufacturing process are shown in Figure 30. This manufacturing process is not summarized in full detail, but details are provided about each of the critical steps. The unidirectional cores were constructed from two components representing the leading and trailing edges of the core. For each of these "halves" of the foil, a sandwich mold, CNC machined, with the foil geometry was capped with ends for carve outs to support subsequent tooling. An additional top mount plate to seals and compacts the composite. The products produced from this molding process are shown in Figure 30a.

At the end of the molding process the joint between the leading edges of the molds is straight. These are subsequently machined to produce a smooth, complex mating surfacing (Figure 30b) between the two. An adhesive is then applied to the mating surface and the foils are held together in custom tooling during the bonding process (Figure 30c). At this point, the foils are longer than intended and the ends are machined down to taper each end of the foil in preparation for mating with the hinges (Figure 30d). At this point, temporary plugs are installed on the end of the foils, the foils are wrapped with four plies of 45° twil, the foil is cured, the plugs are removed, and each end is machined again to align with the taper. Lastly, additional custom tooling supports the process of mating the foils and hinges with the chosen preset pitch angle and maintains the position while the adhesive cures.

A single foil was fabricated for testing and process validation prior to producing a full set of foils. The tests verified that the composite blade did not sustain damage under loads that exceed the

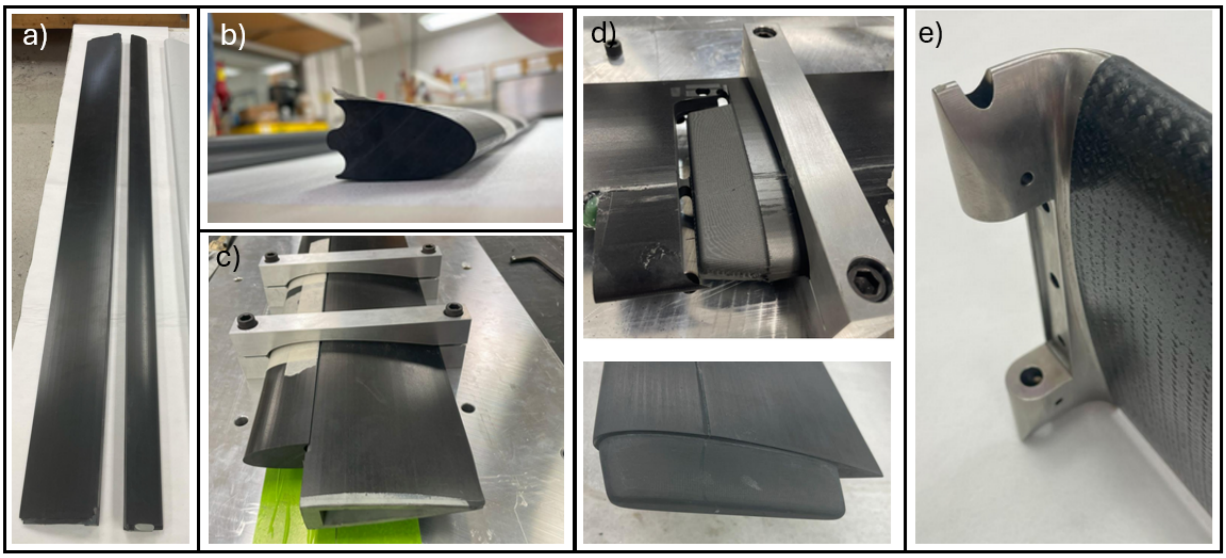

Figure 30: (a) Trailing and leading edges of the uniaxial core after being removed from the molds. (b) A leading leading edge showing the machined surface that mates with the trailing edge. (c) A full foil clamped while the adhesive cures. (d) Machined tapers prior to installing the hinges. (e). Fabricated foil after wrapping the core and installing the hinges.

predicted maximum single-blade loading during an in situ deployment. A test stand consisting of an aluminum I beam with several clamps and a square channel to distribute the point load (Figure 31) was fabricated. When installed, the foil was pinned at one end while the other could translate linearly to accommodate bowing of the foil and the associated deflection of the hinge under load. The fixture was designed to redistribute the force applied at the center of the blade equally to two positions at ⅓ and ⅔ of the blade span.

The test fixture was installed on an MSE Instron 5500R and two tests were carried out on the foil. First, a vertical compression load of 2224 N (500 lbf) was applied to the blade for one minute before increasing in increments of 890 N (200 lbf) up to 9341 N (2100 lbf), holding for one minute at each step. The load was then removed and a second test carried out. Beginning from 9341 N (2100 lbf), the force was increased in 2224 N (500 lbf) increments up to 23130 N (5200 lbf), again holding the load for 1 minute at each condition.

The foil was observed throughout the loading process and several measurements of the deflection of the rotor were made. At the end of testing no damage to the foil (e.g., permanent deformation, cracks) was observed. Under peak loading, the foil deformed approximately 4.8 cm at the center, a value significantly lower than the predicted 7.6 cm derived from analytical estimates. During these tests we estimate that the peak loads exceeded thoses expected in situ by a factor greater than four using conservative load estimates.

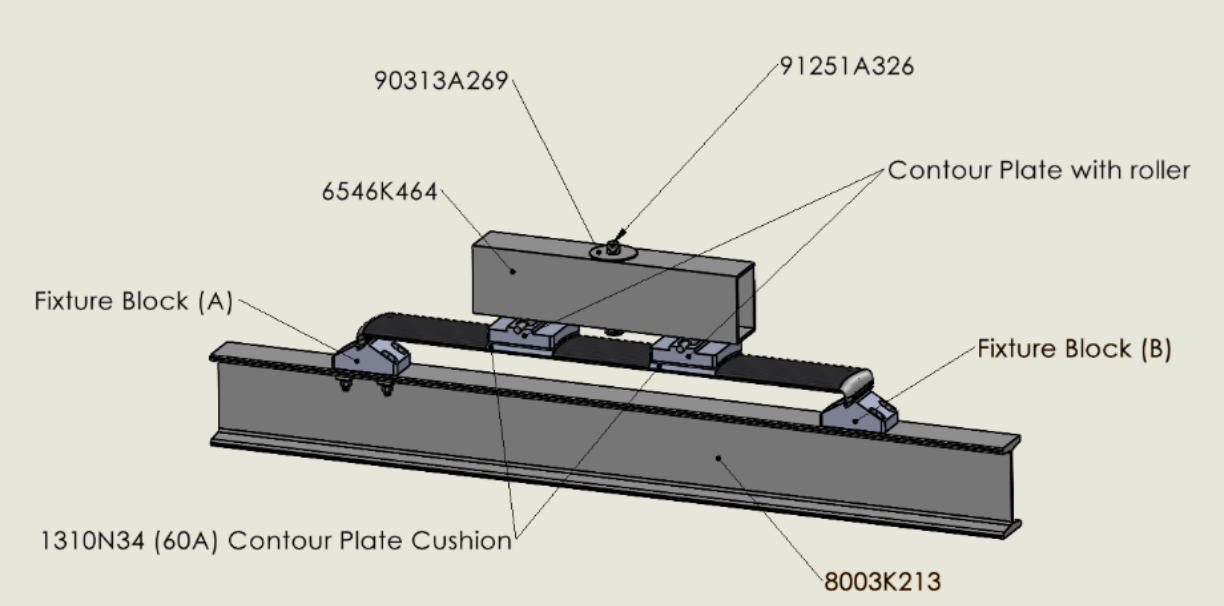

Figure 31: An annotate drawing of the test fixture used to validate the first foil under loads exceeding the maximum in situ load conditions.

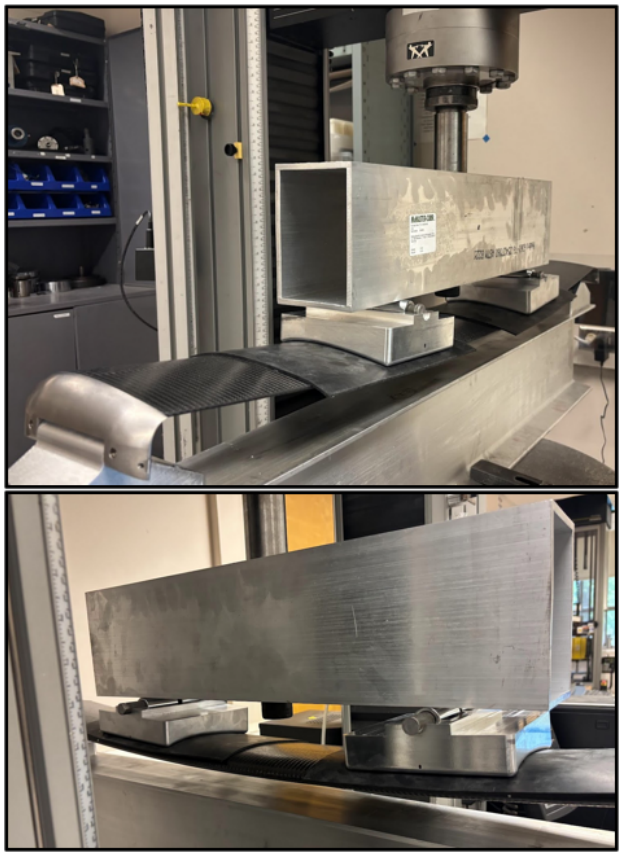

Figure 32: Above and below: the test article (the first foil from production) flexing under load during validation tests.

5 PERFORMANCE ASSESSMENT

This section summarizes the performance of the 2nd generation system including the integration of the MPC, installation of the new rotor, removal of the generator seal, and change of the bearing pack lubricant. Images of the new rotor, endplates, and strut hinge assemblies are shown in Figure 33. Note that compared to the 1st generation design, the rotor appears streamlined and has no protruding parts (e.g., fasteners). The struts and hinges have been further streamlined to lower the drag coefficient of the rotating components. While the overall impact of the modifications is anticipated to be smaller than some of the other modifications made to the system, prior results highlighting the impact of mounting structures on cross-flow turbine performance[13] informed the efforts to streamline the rotor.

In the following sections the power generation values represent a full power budget. That is, only two terms are included: The power drawn from the power supply (if any) and the power dissipated by the load dump (if any). Thus, these values represent power generation capacity of the system. This is a critical consideration in small-scale systems in which even modest hotel loads can consume large portions of the total power budget, thereby limiting system viability.

5.1 Constant Torque Control

Several tests of the 2nd generation design were performed under constant torque control. Under this regime the rotor's speed oscillates as the hydrodynamic torque changes throughout the rotation because the motion is resisted with a constant torque. More complex torque control schemes are often used because they allow operation without inflow measurements so long as the resistive torque applied does not drive the rotor to stall.

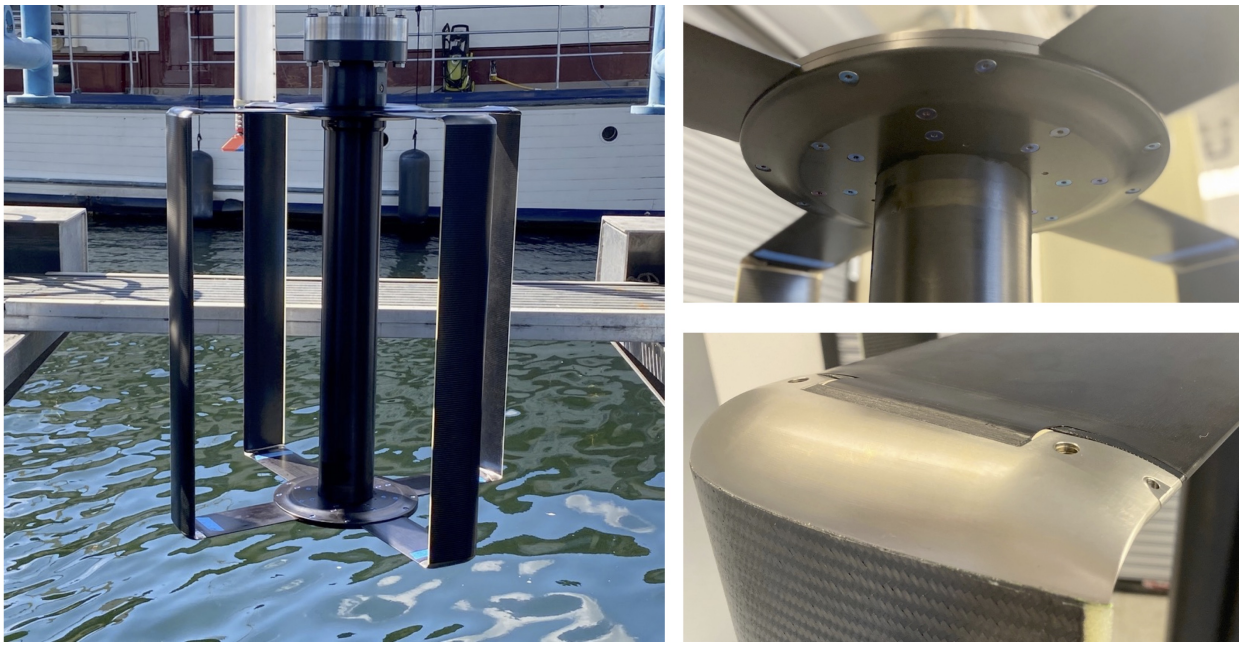

Figure 33: (left) The 2nd generation rotor assembled and installed on *RDL* prior to performance testing. (top right) The upper endplate, shaft, and struts. (bottom right) One hinge/strut assembly.

To test the rotor under constant torque control, *RDL* was operated at a chosen speed and constant heading and the rotor was initialized under constant speed control. Once the targeted speed was achieved, the rotor was shifted to a constant control torque regime and the resistive torque applied by the rotor decreased to 0 N-m. Under these conditions the rotor "freewheels" or spins at its maximum velocity for the inflow speed (this generates no power because no torque is applied). Control torque is then increased in small increments, holding each setpoint for 30 s until the rotor stalls. The torque settings tested ranged from 0.5 N-m to more than 250 N-m, with increments varying from as little as 0.5 N-m at 0.75 m/s inflow to 4 N-m at 2 m/s. Constant torque control tests were not performed at higher speeds because this is primarily implemented above the peak of the tip speed ratio curve (overspeed control). At high inflow speeds (> 2 m/s) significant overspeed control could result in loading the turbine past its design specifications and the rotor spinning at a velocity that causes vibrations in *RDL* equipment. Thus, to protect the test equipment and avoid inadvertently loading the turbine beyond its design specifications, tests at higher velocities were only performed under constant speed control (Figures 35 and 36).

Results from the power generated at vessel speeds of 0.75 – 2.0 m/s are shown in Figure 34. Power generation was measured using the new rotor integrated with the new power electronics after the generator seal was removed and bearing pack lubricant changed, and can be compared to results of similar tests performed with the new power electronics, but with the 1st generation rotor, the generator seal intact, and higher viscosity lubricant in the bearing pack (Figure 15). Performance differences are clear. First, with the old rotor tests could not be performed at 0.75 – 1 m/s because lower torques generated by the rotor coupled with the higher mechanical losses would cause the rotor to stall. The new rotor, however, produces power in constant torque control at 0.75 m/s but all of the power (and more) was consumed by the power electronics. This is, nonetheless, a positive result as it means torque control can be applied at these velocities in situ without stalling the unit; this enables torque control schemes to be implemented across a broader range of conditions. Net power generation in constant torque control was achieved at 1 m/s over a limited range of torques, while at higher inflow speeds positive power generation was achieved widely.

The significance of the changes made to the system are clear in comparisons between ranges of control torques that can be achieved. For the 1st generation system at operating inflow speeds of 1.5 m/s, control torques of approximately 30 N-m would lead to stall, while the 2nd generation system could operate across a range of control torques more than three times as large (roughly 100 N-m). This pattern of extended control torques continues into higher speed operation with control torques above 200 N-m being applied for efficient operation at 2 m/s. This pattern of larger operational ranges emerges as a result of higher torque production by the rotor and reduced parasitic losses in the system and makes better use of the operational capabilities of the generator and the system's magnetic coupling. Furthermore, by trading higher torques for lower speeds, the system achieves higher overall PTO efficiency as recommended by Bassett et al.[2] based on dynamometer characterization of the 1st generation system. By broadening the range of control torques that can be applied at a given setpoint without stalling the rotor, the system will be more capable of operating in turbulent environments without stalling, therefore facilitating the use of adaptive torque control strategies (see Section 5.4).

5.2 Constant Speed Control

Two sets of constant speed control characterization tests were carried out at inflow speeds of 0.75 – 2.5 m/s. These tests, initially carried out at a power converter bus voltage of 277 Vdc were

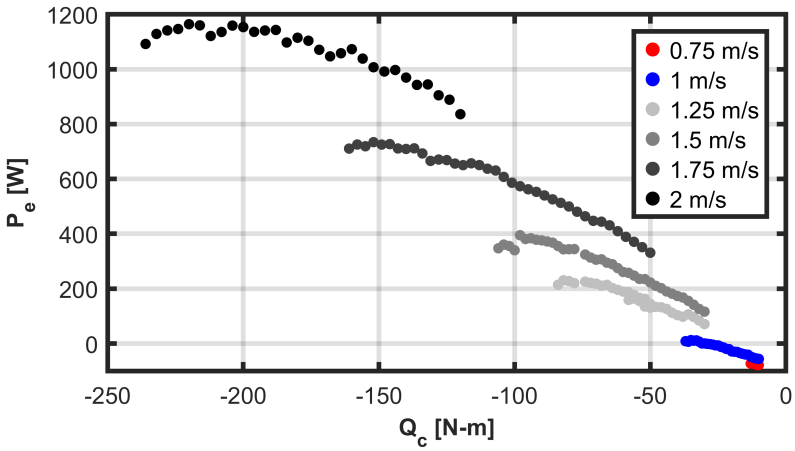

Figure 34: Electrical power output versus control torque from constant control torque tests at six nominal vessel speeds of the 2nd generation system. Tests performed at 0.75 m/s show that although the rotor did not produce net positive electrical power, the rotor could be operated without stalling (the negative power production represents the difference between the power generated by the rotor and power the electronics). Net positive power was produced at all inflow speeds exceeding 1 m/s and the range of control torques that could be applied without stalling was significant larger than could be achieved with the 1st generation rotor.

replicated at 325 Vdc. The initial tests targeted a range of tip speed ratios of $\lambda = 1.4 - 2.4$ in increments of approximately 0.1 based on the assumption that the peak efficiency would occur in the vicinity of $\lambda = 1.9$ as was the case with the 1st generation rotor. However, initial testing showed the the peak shifted to roughly $\lambda = 1.75$ (Figure 35) so the target rotation speeds were shifted down in the second test. In each set of tests the vessel targeted a specific speed over ground while the rotor's speed was stepped through setpoints, each of which was maintained for 60 s. Data were truncated to remove the first and last 5 s of each time series and average values for inflow speed (calculated from aligning timestamps with Nortek Vector ADV measurements), inflow power, and total electrical power generation were calculated. Total electrical power generation across setpoints, presented in both shaft speed and tip speed ratio are included in Figures 35 and 36. Note that in these figures the average inflow speed (not the nominal vessel speed) across each test is noted. Some individual data points are small outliers compared to other points at the same targeted inflow speed. In most cases the variability in vessel speed during a test (e.g., all speed setpoints at 1.5 m/s) was less than 0.05 m/s, but extreme samples did vary by as much as 0.1 m/s. These points were not removed from the data because they do not interfere with the interpretation of broader trends.

The improvements in the 2nd generation system are clear when comparing the results in Figures 35 and 36 to Figure 14 (the old rotor). Whereas the 1st generation rotor only generated net positive power after accounting for power electronics loads at currents exceeding 1.25 m/s, net positive power production occurred at some conditions for inflows less than 1 m/s, while power production occurred over a broad range of rotational speeds at 1.25 m/s. Furthermore, power production at inflow speeds of roughly 1.25 m/s were similar to those at 1.5 m/s for the 1st generation system, despite the fact that the kinetic energy corresponding to the inflow at 1.5 m/s is roughly 72% higher than at 1.25 m/s. Improvements in power generation of roughly 200 W or more are seen across all inflow speeds greater than 1.25 m/s.

Note that there are clear differences between the two sets of speed control tests. These are driven by multiple factors. First, while the nominal vessel speeds were the same for each test, actual vessel speeds were different (e.g., for the 1 m/s tests the actual vessel speeds were 0.94 m/s in one case and 1.06 m/s in the other). This simply reflects challenges in replicating similar vessel speeds under changing environmental conditions with the control systems available on the vessel. Other more subtle differences are associated with the electrical efficiency of the system (see Section 5.3 for details).

While the results of constant speed control testing demonstrate the significant performance improvements with the new rotor, results of the 1st generation Turbine Lander deployment ambiguously demonstrated that *RDL*-based performance tests alone are poor predictors of in situ performance. Although biofouling is a challenge, the complex inflow conditions in energetic environments pose challenges for speed control approaches, which require detailed knowledge of the conditions experienced by the rotor at all times. Thus, while speed control tests are suitable for broad characterization or rotor performance and yield results comparable to those achieved under constant torque control[5], more complex approaches that can adapt to rapidly changing environmental conditions are needed to maximize power extraction during in situ deployments.

5.3 MPC Voltage Bus Impacts on Performance

The new power electronics, which will help the system interface with a battery storage system, can operate over a broader range of direct current voltages. Of relevance to this project are the ranges under which the system has been operated (up to 370 Vdc) and voltages as low as 220 Vdc, which is roughly the lower limit that could be selected without constraining rotor operations at the upper limit of the system's design specifications. The Turbine Lander's generator is a brushless dc motor with a voltage constant of 1520 V/1000 RPM. As described in the following sections, the new rotor operates most efficiently at inflow conditions of 30–90 RPM, which correspond to rms voltages of 45–135 V. The MPC rectifies and steps up these voltages to a target voltage set by the unit, which may be directly programmed or selected/controlled to interface with power storage. To limit the electrical currents and associated losses, the default value of the MPC was 325 Vdc. Thus, during rotor operations, voltages would be stepped up by 190 – 280 V.

Tests were performed on *RDL* under realistic operational conditions to determine how voltage bus settings impact power production. Tests were performed under control using a similar methodology, but restricted to torque ranges near the peak of the performance curve. These tests were performed at inflow speeds of 1 – 2.25 m/s at five different bus voltages: 247, 267, 287, 307, and 325 Vdc. Each condition (control torque, voltage setpoint, inflow speed) was tested for 30 s and the time series was truncated to approximately 25 s under which conditions were stable. The average power output from the system (i.e., power delivered to the load dump) was calculated. These results were plotted for each inflow speed to show the impact of modifications to the system's default voltage.

Our a priori expectations were that dropping the bus voltage would increase power output at low inflow speeds because the lower voltages are associated with the slower rotation rates for peak rotor efficiency. On the other hand, we anticipated that at high inflow speeds losses associated with stepping up the voltage would be lower, but higher electrical currents at the same power production would lead to higher I^2R losses and therefore lower overall power output.

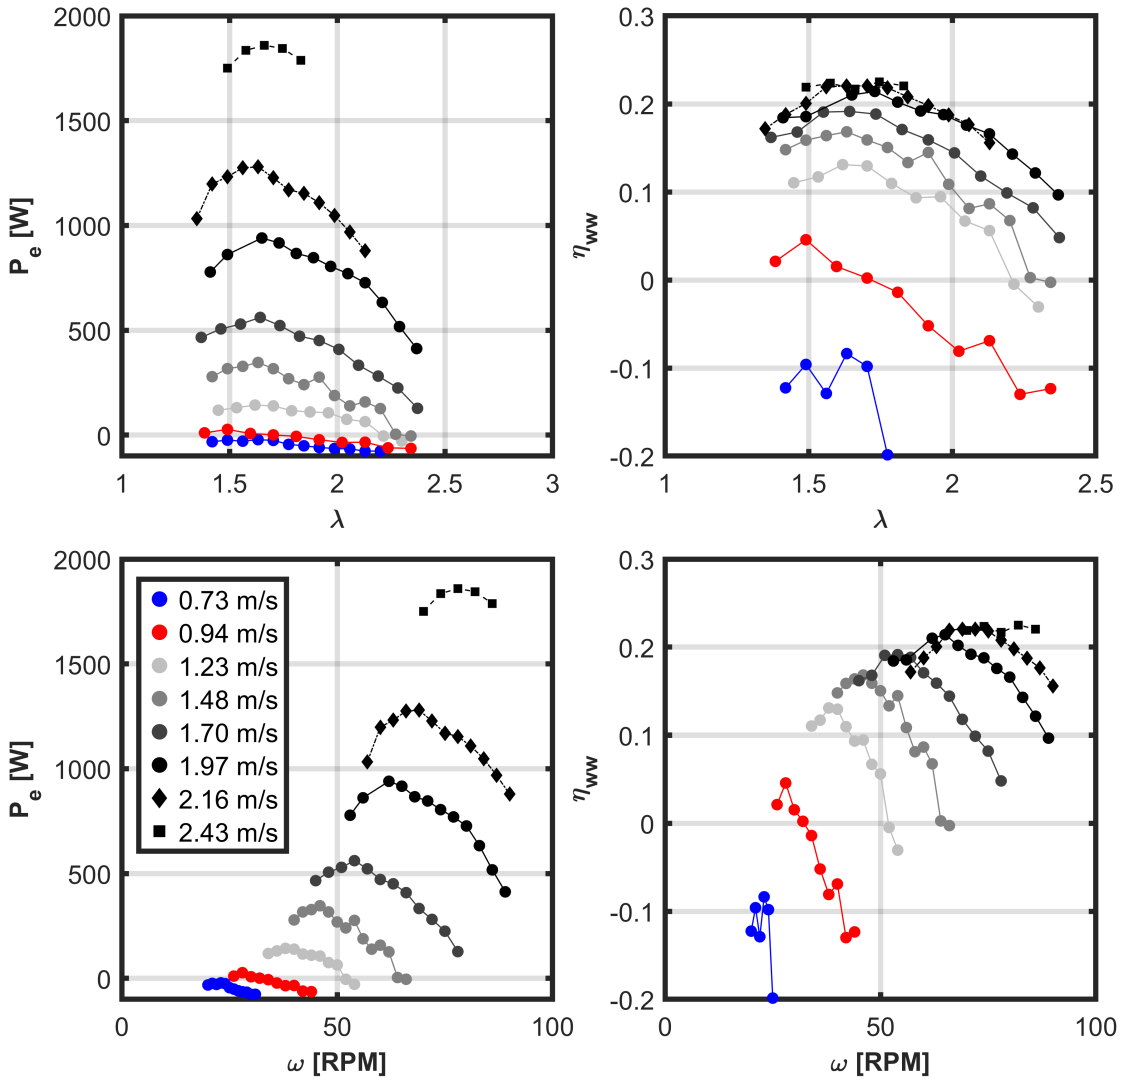

Figure 35: Total power generation (left column) and water-to-wire efficiency (right column) as a function of tip speed ratio (top row) and rotation speed (bottom row) for nominal inflow speeds of 0.75 – 2.5 m/s. Tests were performed with an MPC bus voltage of 277 Vdc.

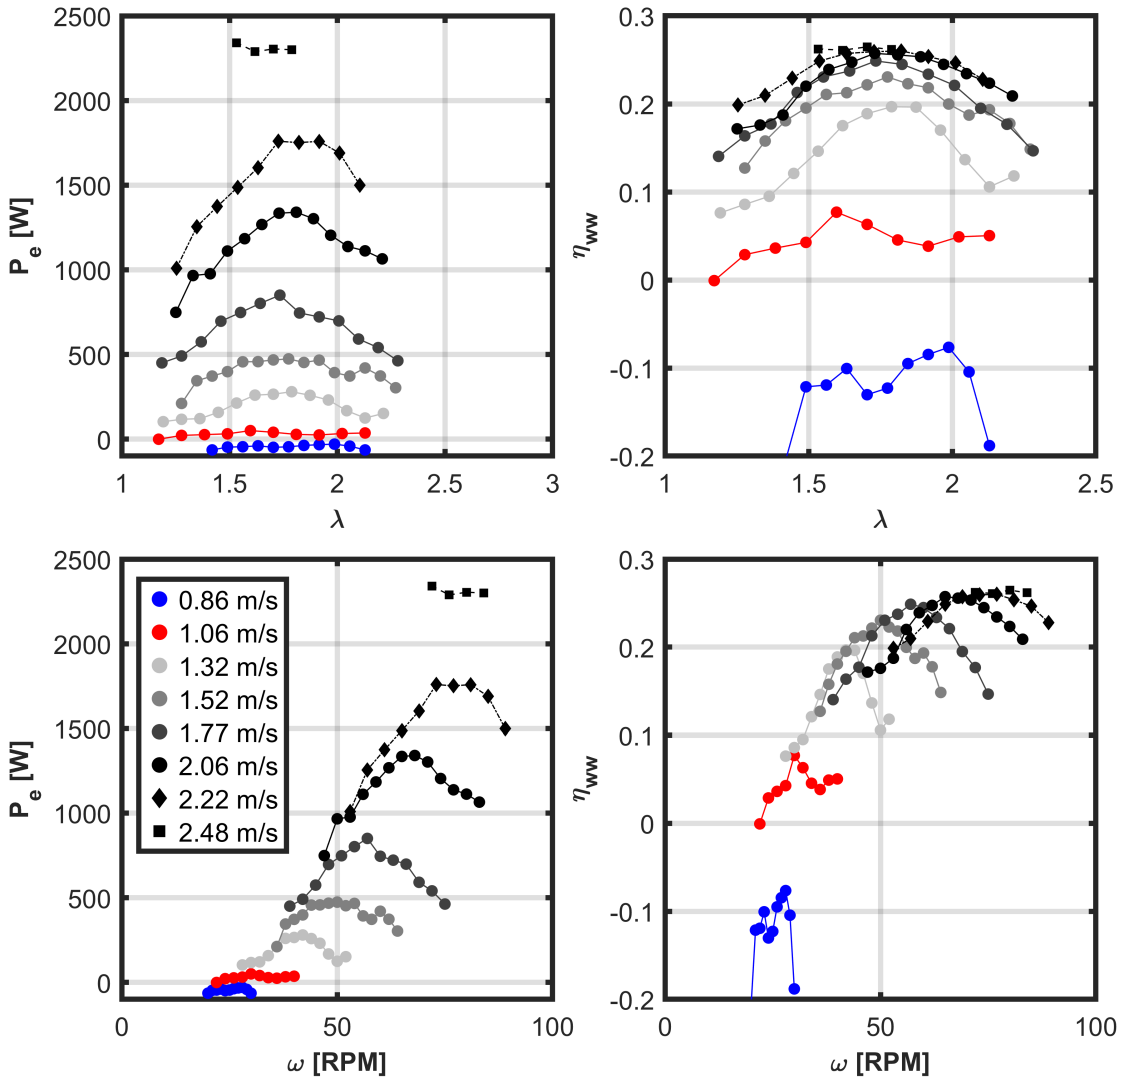

Figure 36: Total power generation (left column) and water-to-wire efficiency (right column) as a function of tip speed ratio (top row) and rotation speed (bottom row) for nominal inflow speeds of 0.75 – 2.5 m/s. Tests were performed with an MPC bus voltage of 325 Vdc.

Table 9: Power output near system cut-in speed. To measure differences in electrical power output the vessel's speed was maintained to the greatest extent possible as the rotor was held at 32 RPM under constant speed control. Inflow speeds were calculated from velocimeter measurements. Despite the fact that the available kinetic power of the water increased slightly as the voltage bus was increased, the system generated the most power at low voltages. We attribute these losses primarily to inefficiencies in boosting the voltage.

V_{bus}	U [m/s]	P_m [W]	P_e [W]	η_{ww}
247	1.04	640	81	0.13
267	1.04	640	64	0.10
307	1.05	660	47	0.07
325	1.09	740	52	0.07

Unfortunately, but perhaps unsurprisingly, the results of these tests proved difficult to interpret due to inherent challenges in maintaining adequate vessel control to create the stable inflow conditions required to quantify subtle differences in overall power generation. Tests were therefore repeated at a single condition to test the hypothesis that differences would be observable. We therefore focused on inflow conditions of 1 m/s, where voltages generated by the rotor are relatively low and overall power production is small. This, in theory, would alow subtle differences to be more easily observed. Furthermore, to simplify the tests we chose a single operating point for the rotor (32 RPM) in constant speed control to minimize differences in loads on the turbine in hopes of maintaining more consistent speeds. Lastly, the length of each test was set to two minutes.

We observed that, as expected, increases in power production could be achieved by lowering the bus voltage at lower inflow speeds (Table 9). These losses could be approximately 30 W at lower power production states. Because power production is relatively low (as is the available kinetic power of the inflow), this value represents an increase in overall power generation. In fact, if the results in Table 9 are replicable this would represent an increase of 50% in net power production at an inflow speed of 1 m/s due to the change in voltage bus. This general principle informed storage system design and will allow for lower voltage operation. Because relatively low energy conditions are more common than periods with strong currents, these modest gains integrated over time can have significant implications for total energy production.

Although the results presented in Table 9 are consistent with the underlying physics and our a priori expectations, the precise gains from operations could not be determined with high precision during testing on *RDL*. To better understand these benefits it would be worth reinstalling the power take-off on the dynamometer to perform controlled testing of the impacts of the voltage bus. Similar voltage bus considerations are required across marine energy converter concepts (including wave energy converters) and better understanding of the combined trade-off could inform future cooperative design efforts.

5.4 Adaptive Torque Control: $K - \omega^2$

A fundamental challenge of speed control in turbine operations is the ability to acquire and process inflow data to control the rotation rate of the turbine such that it is maintained at, or near, the maximum tip speed ratio as the inflow currents change. Given that turbulence resulted in highly variable inflow conditions that could not be adequately tracked, the 1st generation Turbine Lander

in Sequim Bay significantly underperformed relative to predeployment testing.[2] To deal with this challenge, a widely implemented torque control scheme augments the control torque applied by the generator based on the rotor's rotation rate. When implemented well this approach serves as a maximum power point tracking scheme. The control scheme is described by

$$Q_c = K\omega^2, \tag{20}$$

where K is a gain value applied so that the system follows its peak performance curves.[6;16] Following Forbush et al.[6] the gain is determined based on rotor characterization or a system model such that

$$K = \frac{1}{2}\rho A r^3 \frac{C_p^*}{\lambda^{*3}}, \tag{21}$$

where * represents the maximum in the C_P versus λ curve for the system. Because Eqs. 20 and 21 do not rely on the inflow current speed, this approach allows for controls that are independent of inflow conditions. Instead, Equation 21 shows that the relationship between control torque and rotational speed is governed by the rotor's performance.

This approach, referred to here simply as $K - \omega^2$, was implemented in the Turbine Lander deployment in Sequim Bay but resulted in stalls under a broad range of conditions due to overall system performance. One objective of the system redesign was to enable $K - \omega^2$ over a broader range of inflow conditions without stalling. Tests of the $K - \omega^2$ controls with the new rotor were performed on *RDL* to simplify modifications. First, the control torque applied to the rotor was based on a 20-point moving average (0.4 s) of the rotor's rotation. Because the precise C_P^* for the rotor is unknown, K values of 1.5 – 2.7 were tested. Tests of the control scheme were carried out by starting up the system in freewheel with the vessel moving at a speed greater than 1 m/s. The vessel's speed was then increased to approximately 2 m/s and back down to 1 m/s over a period of roughly 5 minutes. Throughout each test the same K value was maintained. In some tests, maximum and minimum velocities were modified to be as high as 2.25 m/s and as low as 0.6 m/s.

In post-processing 5-s moving averages of the inflow conditions were aligned with 5-s moving average outputs including rotor speed, control torque, electrical power production, efficiency, and tip speed ratio. Several tests were combined to provide quantitative comparisons among different K values. For example, data obtained for $K = 1.5$ (Figure 37) shows that over the 15-minute test vessel speeds were 1 – 2 m/s and rotor speeds (during power production) were 40–80 RPM while control torques were –2 – –110 N-m. Maximum power production and water-to-wire efficiencies were approximately 1 kW and 0.2, respectively. These results demonstrate that the control torque and rotor speed respond as expected to the modulating inflow conditions and, most importantly, that the unit did not stall during the test period. This contrasts with the 1st generation rotor, which regularly stalled at inflow speeds less than 1.5 m/s.

In total, eight K values between 1.5 and 2.7 were tested and power production versus inflow speed and η_{ww} versus λ are shown in Figure 38 for four of the tests. Across the tested range of K values the rotor only stalled when inflow speed dropped below approximately 0.8 m/s, which roughly corresponds to system's cut in speed. Across the range of K values we observed that at high inflow speeds the tip speed ratios were above the peak values ($\lambda \sim 1.75$). This represents overspeed control and indicates that the system was operating above its peak efficiency. Thus, it explains why power production under $K - \omega^2$ modestly underperformed testing in constant speed control. As currents decreased, we observed that the same K values did push the system closer towards its optimal λ.

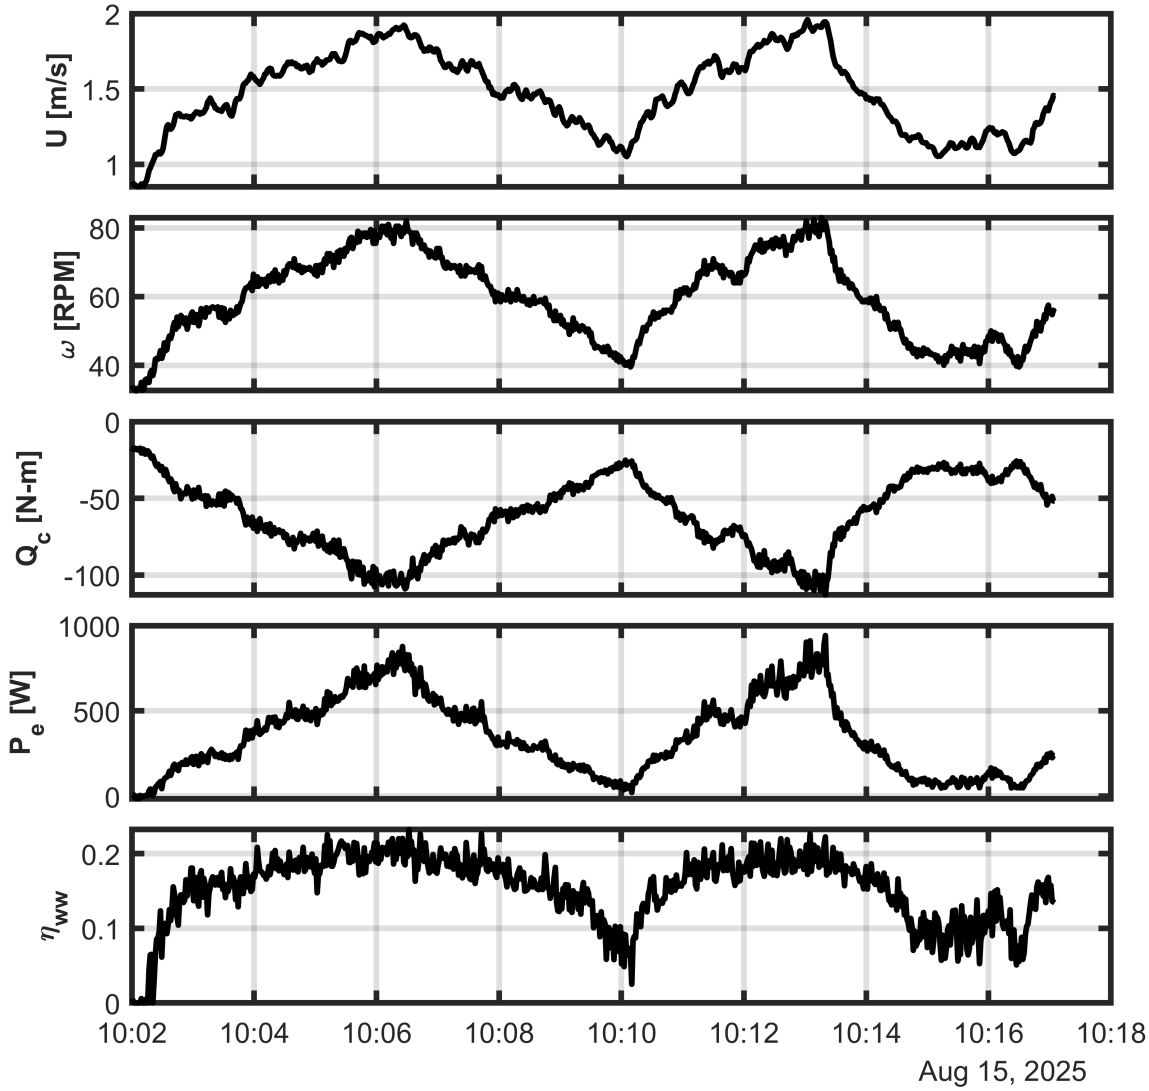

Figure 37: $K-\omega^2$ controls testing using a control gain of $K = 1.5$ (a value well below the predicted optimal value according to Eq. 21). Vessel speed, rotor rotation rate (ω), control torque (Q_c), electrical power output (P_e), and η_{ww} all track the inflow velocity. This and similar examples using different K values suggest that a $K - \omega^2$ controller chould be implemented during in situ testing. Like the other K values tested, $K = 1.5$ resulted in notable overspeed control and therefore increased structural loads and decreased efficiency relative to peak tip speed ratios. However, this overspeed control increases the range in torque between operating and stall conditions, thereby decreasing the risk of stall.

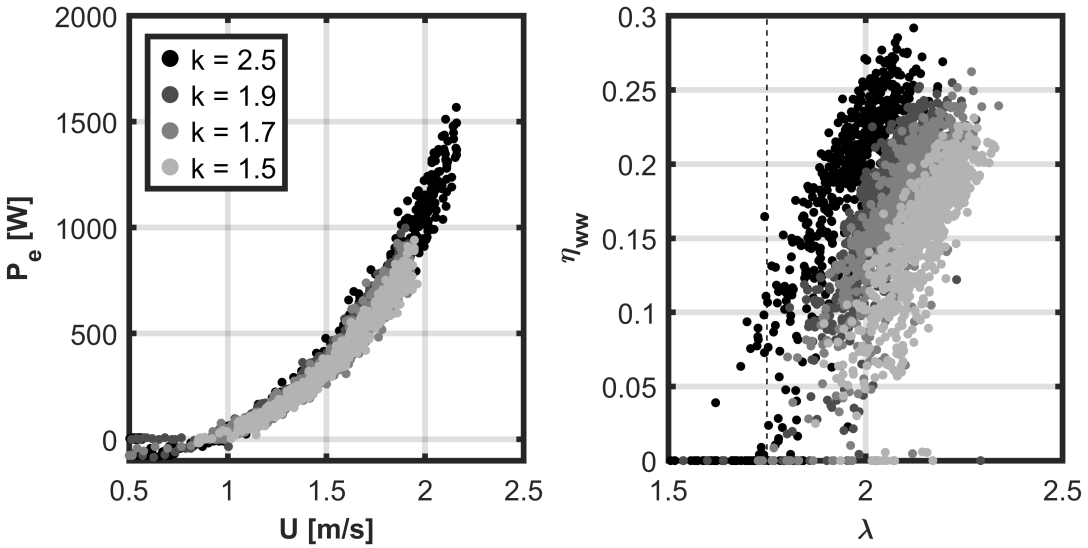

Figure 38: (left) Electrical power generation versus inflow speed for four K values in $K-\omega^2$ control testing using the 2nd generation system. (right) Water-to-wire efficiency versus tip speed ratio for the same four gain settings (inset). The relatively high λ values compared to the peak of the efficiency curves suggests larger gain values would yield modestly higher power production and efficiency values while still operating in overspeed control, thereby minimizing the potential for stall.

Given that the system avoided stall and generated power over a broad range of inflow conditions, it is clear that the modified system is improved and should enable $K-\omega^2$ control to be implemented reliably in future field deployments. To maximize performance we recommend exploring an adaptive $K-\omega^2$ approach that applies a higher K value for periods with stronger currents and a lower K value at lower rotational speeds corresponding to weaker currents. We anticipate that implementing this in situ will require an iterative approach. The goal will be to improve power production by shifting the operating λ closer to the peak ($\lambda \sim 1.75$), while providing enough buffer between the control torque and stall torque. Based on Eq. 21, assuming $C_P = 0.3$, a value of $K \approx 3.5$ may be appropriate for currents greater than 1.5 m/s, while a value closer to the upper limit of testing on *RDL* (e.g., $K = 3$) may be more suitable at lower speeds. These values, however, neglect the impact of vertical shear in the water column and therefore it may be necessary to reduce values further from those predicted by Eq. 21 for stable operations in situ.

Development of an adaptive $K-\omega^2$ controller in situ will require additional work due to the complexities of the inflow and potential for biofouling. Nonetheless, the fact that the 2nd generation design clearly allows for $K-\omega^2$ controls to adapt to the range of inflow conditions under which the system can generate power is promising. Enabling system controls that are independent of inflow measurements is expected to yield significant gains in situ and would make the approach used to control the Turbine Lander more consistent with approaches adopted elsewhere in wind and tidal energy applications.

5.5 Discussion

System improvements include the integration of new power electronics, mechanical modifications to reduce inefficiencies, the new rotor design, and efforts to facilitate integration with energy storage. These provide a clear path towards an autonomous system capable of providing persistent power for applications such as vehicle recharge and advanced sensing at sea. Performance improvements made during this project were only achieved by identifying and addressing several system inefficiencies. Efforts to improve the system are detailed here, providing specific motivation for each modification and insights into the design, fabrication, and characterization trajectory of the project. The broader objectives have been achieved but the path to transitioning a system like the Turbine Lander towards reliable, in situ use is ongoing.

Laboratory and field testing on *RDL* are a reliable way to demonstrate proof of concept and measure system performance, but ongoing, long-term testing in ocean environments is needed. We anticipate that the improvements summarized herein will result in significant increases in net power generation without requiring modifications to the Turbine Lander's foundation. What remains uncertain is how much improvement will be realized in situ. Will the efforts to mitigate biofouling significantly decrease fouling by eelgrass and kelp? Will modified control schemes facilitate significantly better tracking of inflow conditions and prevent stalling due to the enhanced torque generation by the rotor and reduced mechanical losses of the PTO? Answers to these questions can only be addressed through a field testing program. We expect additional deployments to yield new methods to improve the system. Nonetheless, our experiences with the 1st and 2nd generation Turbine Lander systems provide broad support for the development pathway we have taken, which has included laboratory dynamometer tests and controlled field testing on *RDL* prior to advancing to in situ deloyments.

6 RECOMMENDATIONS FOR FUTURE WORK

When the 1st generation system was built the Turbine Lander team was able to leverage laboratory studies, but had limited understanding of the how design choices would impact overall performance. Furthermore, the challenges of co-design are not well discussed in the literature. The interplay between design for electrical, mechanical, and hydrodynamic performance was appreciated, but admittedly not well understood. To close these gaps and understand the elements of the system that could be modified to best improve its performance required us to build and characterize the 1st generation system. As with other complex machines, much can be learned through iteration, particularly when important parts of the parameter space related to design and performance are not widely reported.

Both the 1st and 2nd-generation Turbine Lander rotor designs have four straight foils with symmetric NACA 0018 profiles. This choice was largely driven by evidence from laboratory experiments that this foil profile would be functional, if not optimal. There are several ways that these designs could be improved, either to improve performance, reduce loads, or potentially both, but the best foil parameters to achieve the gains remain the focus of applied research activities. An extensive overview is beyond the scope of this report, although the University of Washington is currently receiving support from the Marine Energy Development (MED) program at NAVFAC EXWC to perform research in these areas.

First, unlike axial-flow turbines, cross-flow turbines undergo significant intracycle lateral and thrust forces driven by the blade count. Fewer blades means that intracycle variability is higher although fewer blades can also achieve higher performance (the difference between two foils versus four is relatively modest).[7] In designing small platforms, intracycle forces are used in estimate overturning moments, so there are benefits to decreasing the intracycle peaks. Review of results in the literature suggests the use of four blades represents a good compromise between performance and loads. As discussed in Bassett et al.[2] the four-bladed rotors also result in lower intracycle torque variability than rotors with fewer blades. This higher variability results in higher instantaneous currents in the motor windings and can lead to significant I^2R losses, displayed as increasing in winding temperature. Tradeoffs are necessary, which points to the importance of co-design and developing an understanding of how mechanical components impact other aspects of total system performance (i.e., even though hydrodynamic performance of a 2-bladed rotor may be modestly better than a 4-bladed rotor, design load fluctuations and electrical loss may make a 2-bladed rotor less favorable). Thus, decisions about design should be made by teams reflecting the broad range of systems integrated in a marine energy converter.

Additional overall system performance is likely to benefit from more complex foil and rotor designs. The use of cambered foils coupled with the rotation of cross-flow turbines can benefit performance by increasing power generation during the power stroke and decreasing drag losses during downstream sweeps. While the overall gains may be modest, perhaps no more than a few percent increases in C_P, the fabrication of cambered foils would not require significant modification to foil fabrication processes and could be implemented without significant modifications to the rotor. As additional research on the subject of cambered blades is published, their use should be considered.

Beyond straight-bladed foils there are several options that could further reduce the intracycle and peak loads experienced by operating turbines. Such concepts include helical and swept blade designs. The performance of these concepts and the associated load reductions is an active area of

research. From a practical perspective, small reductions in peak performance of the rotor (e.g., a reduction of a few percent in the coefficient of performance) would be acceptable if reductions in peak loads could be decreased by a larger factor. In addition to allowing for smaller supporting structures or larger rotors, reductions in instantaneous loads on the system would ultimately result in less wear and tear on other components. In addition, these rotor concepts may have the added benefit of helping to shed biofouling from floating flora that occasionally gets wrapped around foils during in situ operations. Continued research and development of these concepts would be beneficial to the broader community interested in riverine and tidal turbines, particularly for small-scale applications targeting power generation with minimal complexity.

Another recommended area for further research relates to the general shape of cross-flow turbine performance curves. In contrast to low solidity axial-flow turbines, typical cross-flow turbines have a narrow peak in the C_P versus λ curves. As a result, failure to adequately track inflow conditions more negatively impacts power generation for cross-flow turbines than for many axial-flow designs. To reduce these impacts it would be valuable to understand what, if any, design parameters can be modified to generate a C_P curve whose slope, $dC_P/d\lambda$, is lower. In other words, how can cross-flow turbines be designed such that the roll-off in performance away from the optimal tip speed ratio is decreased? We hypothesize that solidity and foil geometries are likely to be the dominant factors here, but there is insufficient published information to inform specific design decisions.

The 2nd generation Turbine Lander represents a significant improvement from the initial concept. Throughout this process, a significant amount of new research has been generated and many lessons have been learned regarding key factors impacting performance and survivability. Given the complex parameter space available for turbine design, system optimization requires a thoughtful, holistic view of the design trade-offs. In other words, system optimization requires a co-operative design approach that includes those familiar with all aspects of the system (i.e., hydrodynamics, manufacturing, mechanical, electrical, software, and operation) participating actively in the design process, basing decisions on detailed knowledge of how individual choices about specific subsystems impact others. This co-operative design approach can only be executed once sufficient information is available. We believe that there is now sufficient technical information available to inform such an effort and recommend desktop co-design efforts focused on small-scale axial- and cross-flow turbines to guide future research and development.

REFERENCES

[1] A. Snortland, K. Van Ness, J. A. Franck, A. Athair, O. Williams, and B. Polagye. Experimental identification of blade-level forces, torque, and pitching moment for cross-flow turbines. *Journal of Fluids and Structures*, 138:104403, 2025. doi: 10.1016/j.jfluidstructs.2025.104403.

[2] C. Bassett, P. Gibbs, H. Wood, R. J. Cavagnaro, B. Cunningham, J. Dosher, J. Joslin, and B. Polagye. Lessons learned from the design and operation of a small-scale cross-flow tidal turbine. *Journal of Ocean Engineering and Marine Energy*, pages 1–21, 2025. doi: 10.1007/s40722-025-00411-y.

[3] E. Cotter, C. Bassett, P. Murphy, M. Scott, A. Runyan, J.M. Almokharrak, L.G. Kao, L.M. Ovall, and S.A. McMillen. Observations of fauna interactions with a small tidal turbine. 2025. Submitted to the PLOS One.

[4] A. Hamilton, F. Cazenave, D. Forbush, R.G. Coe, and G. Bacelli. The MBARI-WEC: a power source for ocean sensing. *J. Ocean Eng. Mar. Energy*, 7:189–200, 2021. doi: 10.1007/s40722-021-00197-9.

[5] B. Polagye, B. Strom, H. Ross, D. Forbush, and R. J. Cavagnaro. Comparison of cross-flow turbine performance under torque-regulated and speed-regulated control. *J. Renew. Sustain. Energy*, 2019. doi: https://doi.org/10.1063/1.5087476.

[6] D. Forbush, R. J. Cavagnaro, J. Donegan, J. McEntee, and B. Polagye. Multi-mode evaluation of power-maximizing cross-flow turbine controllers. *Int. J. Mar. Energy*, 20:80–96, 2017. doi: 10.1016/j.ijome.2017.09.001.

[7] A Hunt, B. Strom, G. Talpey, H. Ross, I. Scherl, S. Brunton, M. Wosnik, and Polagye. B. An experimental evaluation of the interplay between geometry and scale on cross-flow turbine performance. *Renew. Systain. Energy Rev.*, 206:114848, 2024. doi: 10.1016/j.rser.2024.114848.

[8] A. Hunt, C. Stringer, and B. Polagye. Effect of aspect ratio on cross-flow turbine performance. *Journal of Renewable and Sustainable Energy*, 12(5):054501, 2020. doi: 10.1063/5.0016753.

[9] B. Strom, S. Brunton, and B. Polagye. Consequences of preset pitch angle on cross-flow turbine hydrodynamics. In *Proceedings of the 11th European Wave and Tidal Energy Conference*, number 4, pages 08D2–4, 2015.

[10] L. Du, G. Ingram, and R. G. Dominy. Experimental study of the effects of turbine solidity, blade profile, pitch angle, surface roughness, and aspect ratio on the h-darrieus wind turbine self-starting and overall performance. *Energy Science & Engineering*, 7(6):2421–2436, 2019. doi: https://doi.org/10.1002/ese3.430.

[11] A. Snortland, A. Hunt, O. Williams, and B. Polagye. Influence of the downstream blade sweep on cross-flow turbine performance. *Journal of Renewable and Sustainable Energy*, 17 (1):013301, 2025. doi: 10.1063/5.0230563.

[12] P. Bachant, M. Wosnik, B. Gunawan, V. S. Neary, and M. Roh. Experimental study of a reference model vertical-axis cross-flow turbine. *PLoS ONE*, 11(9), 9 2016. doi: 10.1371/journal.pone.0163799.

[13] B. Strom, N. Johnson, and B. Polagye. Impact of blade mounting structures on cross-flow turbine performance. *J. Renew. Sustain. Energy*, 10(3):034504, 6 2018. doi: 10.1063/1.5025322.

[14] R. G. Budynas and J. K. Nisbett. *Shigley's mechanical engineering design.* McGraw-Hill,, Boston :, 8th ed. edition, 2008. URL http://www.loc.gov/catdir/enhancements/fy0701/2006024364-d.html. Previous ed. by Joseph Edward Shigley published under title: Mechanical engineering design.

[15] W.C. Young, R.G. Budynas, and A.M. Sadegh. *Roark's Formulas for Stress and Strain, 8th Edition.* McGraw Hill LLC, 2011. ISBN 9780071742481. URL https://books.google.com/books?id=pF2JAwAAQBAJ.

[16] L. Y. Pao and K. E. Johnson. Control of wind turbines. *IEEE Control systems magazine*, 31 (2):44–62, 2011.