Electromechanical Suspension-based

Energy Harvesting Systems for Railroad Applications

Clément Michel Jean Nagode

Dissertation submitted to the faculty of the Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

in

Mechanical Engineering

Mehdi Ahmadian, Chair

Saied Taheri

Daniel J. Inman

Shashank Priya

John A. Burns

March 19, 2013

Blacksburg, VA

Keywords: Energy Harvesting, Railroad, Suspension, Electromagnetic, Truck, Energy Recovery

This page intentionally left blank.

Electromechanical Suspension-based Energy Harvesting Systems for Railroad Applications

Clément Michel Jean Nagode

Abstract

Currently, in the railroad industry, the lack of electrical sources in freight cars is a problem that has yet to find practical solutions. Although the locomotive generates electricity to power the traction motors and all the equipment required to operate the train, the electrical power cannot, in a practical manner, be carried out along the length of the train, leaving freight cars unpowered. While this has not been a major issue in the past, there is a strong interest in equipping modern cars with a myriad of devices intended to improve safety, operational efficiency, or health monitoring, using devices such as GPS, active RFID tags, and accelerometers. The implementation of such devices, however, is hindered by the unavailability of electricity. Although ideas such as Timken's generator roller bearing or solar panels exist, the railroads have been slow in adopting them for different reasons, including cost, difficulty of implementation, or limited capabilities.

The focus of this research is on the development of vibration-based electromechanical energy harvesting systems that would provide electrical power in a freight car. With size and shape similar to conventional shock absorbers, these devices are designed to be placed in parallel with the suspension elements, possibly inside the coil spring, thereby maximizing unutilized space. When the train is in motion, the suspension will accommodate the imperfections of the track, and its relative velocity is used as the input for the harvester, which converts the mechanical energy to useful electrical energy.

Beyond developing energy harvesters for freight railcar primary suspensions, this study explores track wayside and miniature systems that can be deployed for applications other than railcars. The trackside systems can be used in places where electrical energy is not readily available, but where, however, there is a need for it. The miniature systems are useful for applications such as bicycle energy.

Beyond the design and development of the harvesters, an extensive amount of laboratory testing was conducted to evaluate both the amount of electrical power that can be obtained and the reliability of the components when subjected to repeated vibration cycles. Laboratory tests, totaling more than two million cycles, proved that all the components of the harvester can satisfactorily survive the conditions to which they are subjected in the field. The test results also indicate that the harvesters are capable of generating up to 50 Watts at 22 V_{rms}, using a 10-Ohm resistor with sine wave inputs, and over 30 Watts at peak with replicated suspension displacements, making them suitable to directly power onboard instruments or to trickle charge a battery.

This page intentionally left blank.

Acknowledgements

I would like to first thank my beloved and wonderful wife, Audrey, for her unconditional love and support and for making my life more complete and easier throughout that grand adventure that brought us to another continent. I would also like to thank all my family who have always been a great source of motivation and encouragement. I would never be where I am now without them.

I am deeply grateful to Dr. Mehdi Ahmadian for first receiving me as a visiting scholar some six years ago, and then for offering me the opportunity to come back for my graduate studies. He is certainly a great advisor with whom it is a real pleasure to work. He made all this adventure possible and immensely changed, for the better, the path of my career.

I would like to thank Dr. Saied Taheri for his advice and helpful input throughout the project. I would like to extend my thanks to Dr. Daniel Inman, Dr. Shashank Priya, and Dr. John Burns for serving on my committee.

I am really thankful to Michael Craft, CVeSS senior researcher, for his friendship and his invaluable help on probably every single aspect of my project. I am also very thankful to all the members of the CVeSS family: Dr. Alireza Farjoud, William Burke, Khaled Adjerid, Andrew Peterson, Tarek Shalaby, Scott Naranjo, Alex Keylin, and all the others who make life in the lab so enjoyable.

I would like to also thank Benjamin Dupas for lending me a hand (almost literally!) and helping me with some testing. Thanks to James Shultz for the help with the work on the generators.

Finally, I would like to acknowledge the American Association of Railroads for supporting this study and for their help through the project, with a special mention to the people from the Transportation Technology Center, Inc.

And thanks to everyone I may have, shamefully, forgotten!

This page intentionally left blank.

Contents

List of Figures

List of Tables

This page intentionally left blank.

Chapter 1

Introduction

1.1 Motivation

As interest grows within the railroad industry to implement more electronic devices on rail cars, the need for a reliable source of electrical power is more pressing. Although the locomotives generate power to supply the traction motors and all the equipment, the remainder of the train is not electrified. This means that in order to install a monitoring system, a GPS tracking device, or an active RFID tag on a freight car, there must also be solutions for powering it. Often, these electronic devices are simply powered by batteries, limiting the range of possible applications and incurring the risk of losing power. There are a few options currently available, such as Timken's generator roller bearing or solar panels, but railroad companies have been slow in adopting them for various reasons, including cost, difficulty of implementation, or limited capabilities. Therefore, this leaves a real need in the railroad industry for a system that can provide enough power in every freight car to supply electrical devices.

The research work presented in this document is motivated by this actual industry need and seeks to find economically and technically feasible solutions.

1.2 Objectives

The primary objectives of this research can be summarized as follows:

- Investigate the possibility of using linear generators to harvest energy from a freight car suspension
- Design, develop, and test an efficient energy harvesting system using a rotational generator
- Develop a numerical model of the energy harvesting system that can be used as a design or optimization tool, or as an element in the model of a complete car suspension
- Research alternate ways to provide electrical power in freight cars
- Research alternate applications for an energy harvester using linear motion as input
- Study the possibility of miniaturizing the system for implementation in a space-limited environment

1.3 Approach

Prior to designing a system to harvest energy from a rail car suspension, the first step is undoubtedly the research of the systems already available, commercially or filed in patent. It is extremely important to have a thorough knowledge of existing solutions to the lack of electric sources in freight cars and to understand their limitations and the reasons why they have not been fully adopted. It is also essential to research systems that are not designed for the railroad industry, but that could still provide a viable solution, after some redesign.

Different approaches will be taken for each design. The first concept is to use a linear generator to harvest energy from the motion of the suspension. The goal here is to design the generator itself. The approach taken is to simulate the behavior of the magnetic field in a multitude of potential configurations, and then analyze the data to determine which concept is the most promising. The next phase is to build a prototype and determine its performance. Finally, an important aspect is determining whether the simulation matches the results from the prototype.

The second concept is to design a system using a rotating generator. In this case, the goal is no longer to design the generator itself, but to design around it. Because there are already many options for well-engineered generators, it is unnecessary to design a new one; however, knowledge of the generator is extremely useful. The challenge is then to create the appropriate arrangement of parts to obtain an efficient design. The final concept is then modified to fit various applications.

1.4 Outline

Chapter 2 will give an overview of the systems already in existence that could provide some insight into the development of an energy harvesting system for a freight car suspension. Maxwell's equations will be presented, as it is important to understand them in order to design electro-mechanical devices.

In Chapter 3, the work that has been achieved concerning the linear generators will be presented. Because the motion of the suspension is a translation, a linear generator without any mechanical transformation should be the most efficient way to proceed. The analysis of the generator will be presented, as well as the prototypes that were built and tested.

Chapter 4 will introduce harvesting systems that transform the linear motion of the suspension into rotation in order to use rotating generators, recognizing that a gearbox, while introducing some energy losses, offers far better conditions for the generator to produce electricity. The work presented will range from the initial design to the implementation of pre-production systems in the field, with intensive lab testing in between.

Chapter 5 will present the modeling and simulation work that has been achieved to estimate the output voltage and the potential harvested power.

In Chapter 6, the possibility of designing a miniature harvester around the same concept used for a freight car suspension will be investigated. As a demonstration, a mountain bike fork will be used.

Chapter 7 will present the work done to develop an energy harvester for a rail car suspension using a rack and pinion instead of a ball screw. It will discuss the various challenges this concept presents, as well as its potential advantages.

In Chapter 8, a new approach will be presented. In this design, instead of using the motion of the suspension, the system uses the rotation of the axle. An unconventional concept, designed to be clamped around the axle, will be introduced.

Chapter 9 will step away from onboard applications, presenting the work done to develop an energy harvester for wayside railroad applications. The vertical deflection of the track can be used directly with the suspension system, but discussions on how to improve the performance for that particular application will also be included.

1.5 Contributions

The contributions of this research work can be summarized as follows:

- Development of a method to determine the open voltage of a linear generator using FEMM, free finite-element quasi-static magnetic analysis software
- Design of linear generators that can produce useable voltages with small, low frequency inputs
- Design of an energy harvester using a rotating generator, after transformation of the linear motion of the suspension into a rotation

- Development of an innovative gearing mechanism that allows the generator to always rotate in the same direction, therefore improving efficiency
- Development of a system to conduct durability testing on the prototypes
- Assembly of prototypes for field testing and implementation of the harvesters on actual freight cars
- Development of a model for the energy harvesting system
- Design of a system that, based on the same concept, can be fitted in applications with limited available space
- Design of an energy harvesting system for freight car suspension using a rack and pinion
- Development of a generator that can easily be put in place on the car axle and can provide large amounts of power
- Investigation into using the suspension harvesters for providing power to wayside railroad equipment using the vertical deflection of the track

Chapter 2

Literature Review

2.1 Existing Freight Car Powering Systems

Many electrical systems, such as those for communication or monitoring, are mounted on the locomotives, which makes supplying power a non-issue. However, when equipment is mounted on a freight car, having an available, reliable source of electricity may become challenging. As a result, railroads have adopted different work-around solutions.

2.1.1 Batteries

A simple solution to power any electrical device is to use batteries. Some systems have reportedly run on rechargeable batteries, such as a basic lead-acid battery [1]. Others, like RFID tags, can use non-rechargeable types, such as alkaline AA batteries. The use of small batteries is made possible by the low power consumption of the device. A few years of service may be expected, but at the cost of limited capabilities like low sampling rate or turning the transmitter on only for short periods. The main drawback of batteries is therefore the trade-off between life expectancy and performance.

The primary constraint of batteries is that they must be monitored, and when their voltage becomes too low, a system for replacement or recharging has to be in place.

However, batteries are advantageous in that they are readily available, easy to install, and can potentially fit any application.

2.1.2 Timken Generator Bearing

The Timken Company, a world leader in bearing manufacturing, is one of the providers of journal bearings for American railroads. In the 1990s, they developed a special bearing that, between the two

rows of needles, incorporates a system able to generate some electric power [2]. A series of magnets is fixed on the axle and coils are attached to the outer ring of the bearing (Figure 2.1). As the train rolls, the axle rotates, creating a relative motion between the two which in turns generates electrical energy. The power generated can then be used to supply onboard systems directly or to recharge a battery. At 50 miles per hour, the system is able to produce up to 24 Watts [3].

Figure 2.1: Timken generator bearing [4]

In the early 2000s, Timken introduced a similar concept: the Guardian™ bearing, which has a comparable arrangement of magnets and coils between the two needle bearings. However, the power generated is used to power the embedded sensors and transmitter [5]. The system is able to detect temperature, vibrations and rotational speed, and can transmit the data wirelessly. It can send an alarm, for instance, in the case of a bearing or wheel failure or a stuck hand brake.

Those Timken bearings with embedded generators have been tested by the Transportation Technology Center, Inc. in Pueblo, Co and in coal revenue service [6], but they have not yet been widely accepted by railroads, in part due to resistance to the introduction of a non-standard bearing in their inventory.

2.1.3 Journal Bearing Cap Generator

Along with ideas similar to those of the Timken Company to use the rotation of the axle as input power, numerous patents [7,8,9,10,11,12,13,14] have been filed for systems fixed in place of the journal bearing

cap. Typically, a cap is screwed to the end of the axle, maintaining the bearing in position. It can be modified in such a way that it still holds the bearing, and in addition, can drive a generator. When the train moves, the rotation of the axle is transmitted to the rotor of the generator, producing electricity.

There are two different variations to this concept. In one, the generator is placed directly at the end of the axle. The rotor is fixed on the bearing cap or is designed to completely replace it. The rotor will then rotate with the wheel, at the exact same speed. The stator is connected to a fixed element, like the outer ring of the bearing or the side frame of the bogie.

In a second concept, a drive system is used. A sprocket or pulley is attached to the axle, and the generator is driven by a chain or a belt. Often, the generator support is pinned to allow tension in the belt/chain. This system adds some complexity, but it can increase the slow speed of rotation of the axle to a level more suitable for efficient electricity production. It also allows more flexibility in the positioning of the generator, preventing it from sticking out to the side, and potentially allowing the use of a larger unit. Such a system can be seen on the caboose Soo 31 (Figure 2.2) of the Minnesota Transportation Museum. A caboose is a rail car hauled at the end of a freight train to house crewmen. This Soo 31 was purchased by the Soo Line in 1967. The generator primarily provided power to equipment such as a radio to communicate with the locomotive, and light which allowed night-time work.

Figure 2.2: Caboose with axle generator - Minnesota Transportation Museum

2.1.4 Axle Generator

Another way to use the rotation of the axle to produce electricity is to take the power directly from the axle itself, between the wheels. This idea has been in existence for a long time, as shown by patents filed in the 19th century [15]. More patents followed in the 20th century [16,17,18] but the basic concept has remained relatively similar: a split pulley is clamped onto the axle and belts transmit the motion to the generator shaft (Figure 2.3). Various systems are used to put tension in those belts. Special care must be taken in the manufacturing of the pulley, as it can be difficult to balance and may induce vibrations on this critical element that an axle is.

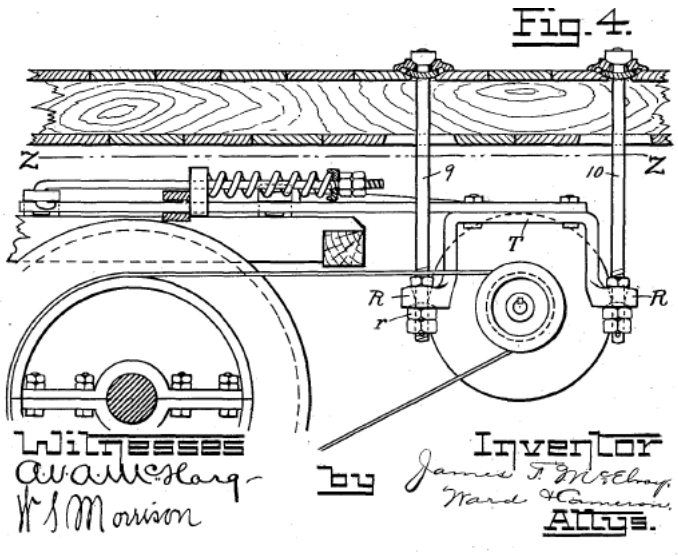

Figure 2.3: Patent drawing - Device for attaching dynamos to railway cars [16]

2.1.5 Solar Panels

Solar panels are not a new technology. However, in recent years, they have become more efficient and more affordable, making them more common. Railroads have used them to power wayside equipment, such as rail lubricators, and they also appear in onboard applications. A team from the University of Nebraska-Lincoln used solar panels to charge, during the day, the batteries of their real-time vertical track deflection measurement system [1] (Figure 2.4).

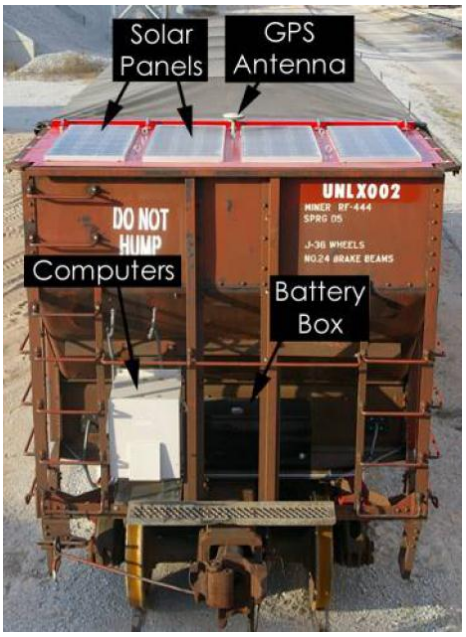

Figure 2.4: Nebraska-Lincoln's track monitoring car

The main drawback of solar panels is that they require sunlight, with a good exposition, to perform efficiently and reach their power rating. At night, on cloudy days, with grazing lights, or in forest areas, the output power may become low or nonexistent. Another consideration is that railroad conditions may be too harsh for solar panels, as they are quite fragile. In addition, although it may be easy to install them on the roof of box cars, it may be impossible to place them on other revenue service cars like car hoppers, for instance. An unexpected challenge faced by railroad companies with solar panels is that they often get stolen or vandalized.

2.2 Linear Generators

A linear generator is a type of generator that converts mechanical energy to electrical energy but, unlike "classic" generators which use a rotary motion, it works with linear displacements. The underlying physical principle remains the same: electromagnetic induction. A magnet moves relative to a coil, changing the magnetic field inside that coil, which induces a voltage. Numerous applications, sometimes leading to patents [19,20], have been derived from this simple principle.

2.2.1 Faraday/Shake Flashlight

One of the most well-known commercially available systems using a linear generator is the shake flashlight, also called the Faraday flashlight, according to the name of a 2004 patent [21]. This torch does not run on batteries, but instead has a built-in linear generator that charges a capacitor. When the flashlight is shaken lengthwise, a magnet slides back and forth inside the handle through a coil of wire. This motion creates a variable magnetic field inside the coil and, as stated in Faraday's law of induction (described in detail in Section 2.5), this in turn generates a voltage. Generally, that voltage is rectified through a diode bridge and energy is stored inside a capacitor or, more rarely, inside a rechargeable battery. A switch then allows the energy to be used in one or more high intensity LEDs.

Figure 2.5: Shake flashlight [22]

However, there are only magnetically reluctant materials (plastic or air) around the coil. This makes a poor magnetic circuit, leading to low efficiency. Therefore, a minimal amount of the physical energy spent shaking the flashlight is actually transferred to the capacitor. Beyond that, these flashlights work well, despite the fact that numerous cheap and poorly-made versions have led to a bad image.

2.2.2 Wave Power Plant

An active area of research for linear generators is in wave energy plants. The idea is to use the motion of oceanic waves to produce electricity. The concept is relatively simple: a buoyant object moves up and down with the waves, and the motion is transferred to a generator fixed at the bottom of the sea. There are

two main designs [23]. The first one uses a buoy floating at the surface of the water (Figure 2.6 – left). The second one, called the Archimedes Wave Swing (AWS), uses a fully submerged gas-filled vessel, which expands and contracts as the waves pass above, thus changing the height of the column of water above (Figure 2.6 – right). The main advantage of this configuration is that because it is totally submerged, the system is better protected against bad weather. The main drawback is that technical constraints limit the stroke, making the system only suitable for relatively long wavelengths, such as the swell from the ocean. The AWS is protected by a patent [24] and is developed by a European cooperation between five companies, three universities, and two research institutes [23].

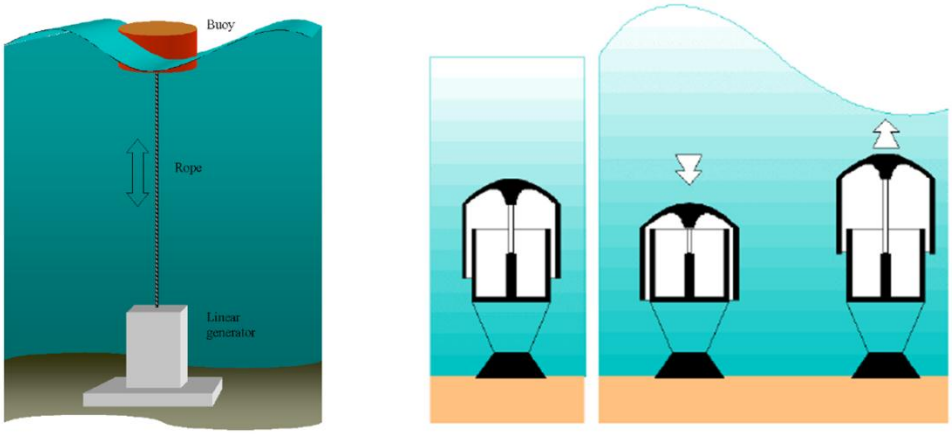

Figure 2.6: Wave energy plants [23]

The concept of electricity production from wave movement gained serious interest after the 1970's oil crisis, when initiatives for the development of renewable energy technologies sprung to life. Technical difficulties and a lack of economically viable solutions caused interest to decline, but the 1990s saw a resurgence of the idea. However, at the time, power stations were only interesting for isolated installations, and their usability in supplying power to a grid had yet to be proven. Until that time, nearly all the wave power plants were using "classical" high speed rotary generators for the mechanical-to-electrical conversion, thus, all of them needed to overcome the difficult task of converting the up and down motion of the wave to a rotation. Several creative solutions have since been developed, but their complexity usually leads to expensive construction and maintenance [23].

To avoid this issue, rotary generators are being replaced by linear generators. The reciprocating vertical motion of the buoy can then be directly coupled to the generator, eliminating the need for a complex

mechanical system. In this configuration, a steel cable is attached to the buoy on one end, and to the rotor of the generator on the other end. The rotor in a linear generator does not actually rotate but is named by analogy to a rotary generator. Shaped like a piston, it is composed of an arrangement of steel parts and permanent magnets. A spring fixed between the ground and the rotor maintains the tension in the cable (Figure 2.7). This type of generator is used on the Archimedes Wave System deployed as a prototype on the coast of Portugal in 2001-2004.

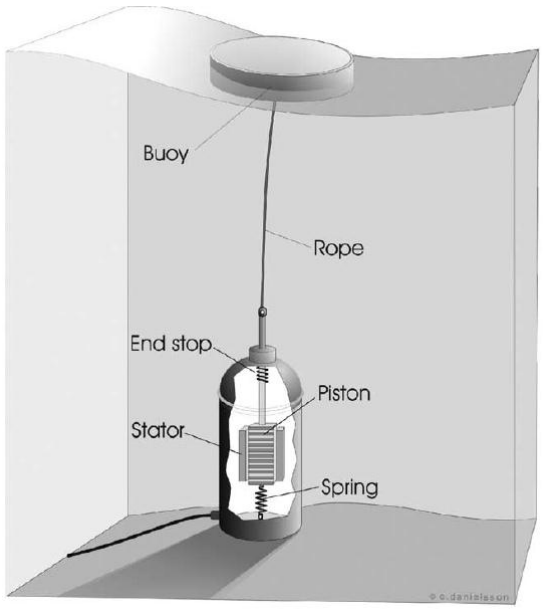

Figure 2.7: Schematic of the wave energy converter with linear generator [25]

Fueled by the trend of green energy, the field of wave power plants has become the source of numerous projects working to improve and further develop linear generators [23,25,26,27,28]. Producing kilowatts of power, they are in an entirely different category from a suspension energy harvester. However, the initial concept remains similar. There is much insight to gain from the design of such a large generator, and once scaled down, the concept can give some guidelines for the development of a linear generator.

Vibration Tuned Harvester

At the other end of the power spectrum from the wave generators, electromagnetic vibration harvesting generators are able to produce few milliwatts to power wireless sensor nodes. As their name indicates, these small devices convert mechanical energy, in the form of vibrations, to electrical energy. A coil of wire is fixed on a vibrating element, such as an electric motor or an engine, and the vibrations cause a magnet to move back and forth inside the coil (Figure 2.8). This motion produces a variable magnetic field, which in turn induces a voltage, as stated in Faraday's law (see Section 2.5).

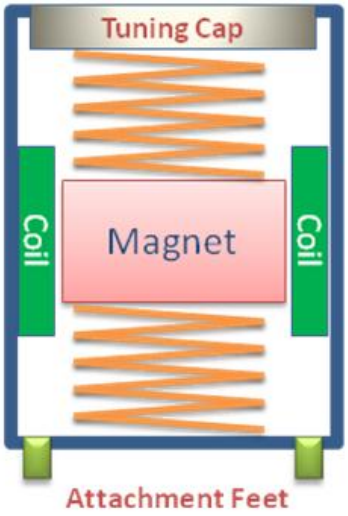

Figure 2.8: Vibration harvesting system

To improve efficiency, the magnet is mounted on springs and the system is used at its resonance frequency. The natural frequency of the system is simply defined by:

$$f_0 = \sqrt{k/m} \tag{2.1}$$

where k is the stiffness of the (equivalent) spring and m is the mass of the magnet.

The natural frequency is chosen to correspond to the excitation frequency (or a harmonic), causing resonance. Thus, the amplitude of motion of the magnet far exceeds that of the input vibrations. This is the exact same principle used in vibration tuned absorbers. Like those vibration absorbers, the harvesting system can only work around that frequency. This is not an issue for applications with motors or generators rotating at constant speed, but it cannot be suitable for frequency-variable applications.

Starting as a spin-off of the University of Southampton, UK in 2004, Perpetuum [29] was the first worldwide to produce and commercialize practical vibration harvesting systems capable of reliably providing the power required by autonomous wireless sensor nodes. Today, Perpetuum is among the global leaders in vibration-based generators, and is now developing a series of micro-generators specifically designed for railroad applications. One designated application is the monitoring of rail car wheel bearings. The system can detect abnormalities and send information wirelessly, so problems can be addressed before parts become damaged and accidents occur. A second targeted application is tracking, so that the location of expensive or hazardous cargo is known at all times. Optimized for the vibration frequencies typical on a rail car (between 30Hz and 60Hz), this system proved during tests on actual rail cars that it can generate 10 to 15mW of continuous power when the train is moving.

2.3 Generative Shock Absorbers

The automotive industry has also started to show interest in harvesting energy from the motion of the suspension. One of the most publicized examples is the patented *Genshock* system [30] developed by a team from MIT. Unlike conventional hydraulic dampers which dissipate energy by passing a fluid, generally oil, through small orifices, thus generating heat, the *Genshock* system reroutes the fluid to actuate a hydraulic turbine that drives an electric generator (Figure 2.9).

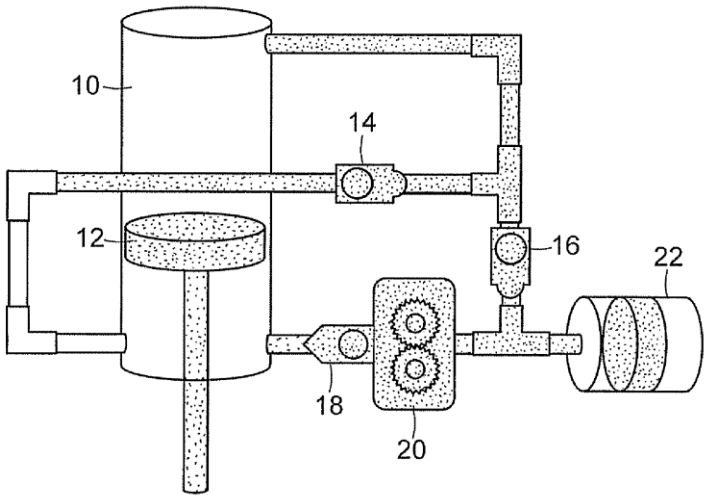

Figure 2.9: *Genshock* system schematics [30]

Although no actual data can be found, Levant Power, the company founded by the group from MIT, states on their website that the system can produce from tens of watts to several kilowatts [31] depending on the vehicle and road conditions, and it could generate up to an average of 1kW when mounted on a heavy truck driving on a standard road. The company is reportedly interested in developing its concept for truck, bus, industrial, marine, and rail applications.

Car manufacturers are also investigating the concept of regenerative dampers. Honda Motor Company filed a patent for an electric damper [32] in which a rotating generator is driven by a rack and pinion arrangement (Figure 2.10).

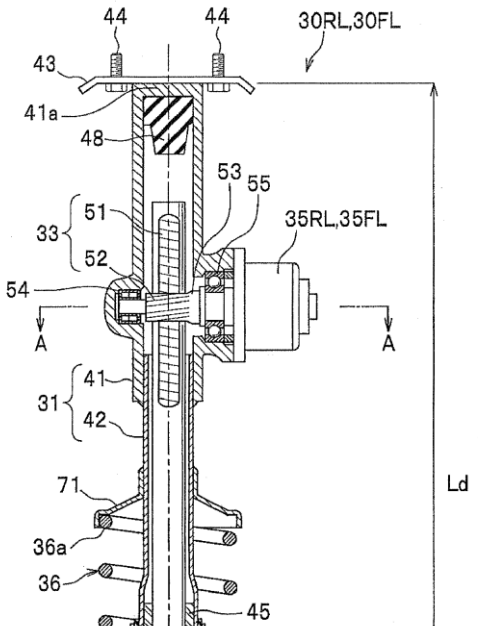

Figure 2.10: Honda's electric damper

Numerous other examples of electric dampers can be found in literature and patents, illustrating the interest in that technology.

The team of Dr. Zuo at Stony Brook University has recently developed an energy-harvesting shock absorber (Figure 2.11) [33]. Their prototype uses a rack and pinion to transform the linear motion of the suspension into a rotation that is then redirected to a gearbox by bevel gears. They improved the design by including one-way bearings and by cleverly using three bevel gears in total to “rectify” the motion and force the generator to spin in only one direction. This design is now protected by a patent [34], published in January 2013.

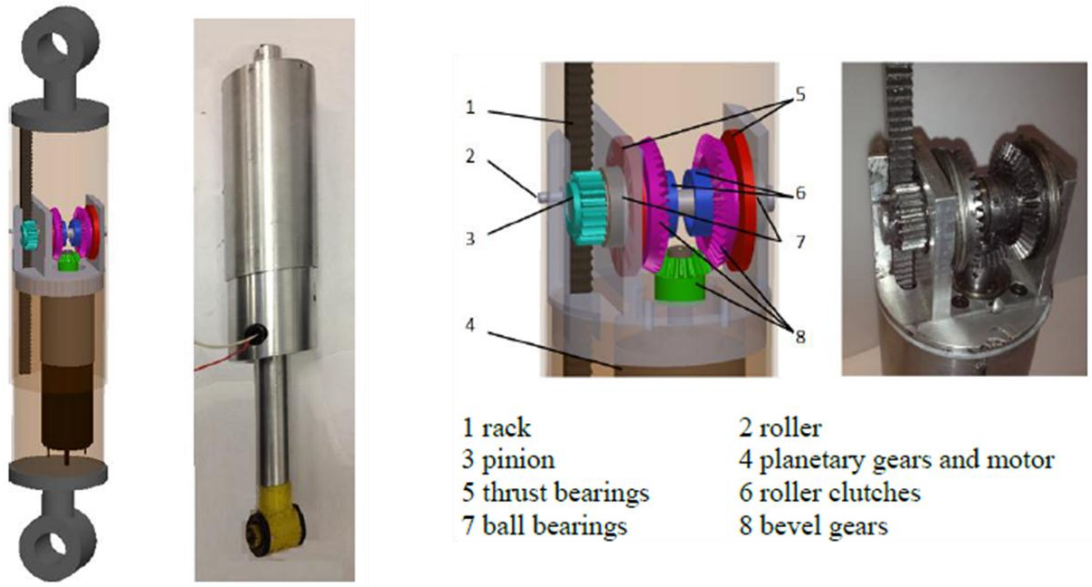

Figure 2.11: Energy-harvesting "shock absorber" from Stony Brook [33]

Linear motors are also used to act as a damper. One of the well-known examples is the fully active suspension system developed by Bose. Using their expertise in voice-coil speakers, they developed an electro-magnetic shock absorber capable of greatly improving the dynamics of the vehicle. Other less famous examples can also be found [35], using permanent magnets and coils to generate force based on Lorentz force law.

2.4 Inerters

An inerter (not to be confused with an inverter) is a mechanical device with two attachment points in which the force is proportional to the relative acceleration between the two ends. The force can be expressed by the equation:

$$F = b\left(\dot{v}_1 - \dot{v}_2\right) \tag{2.2}$$

where b is the inertance, with SI units in kilograms.

By comparison, in a spring, the force between the two ends is proportional to the displacement and in a damper, the force is proportional to the relative velocity. The inerters then complete the set with their

force-acceleration relationship. Inerters were introduced in 2002 by Malcolm Smith, professor at the University of Cambridge, UK. Working on the analogy between electrical and mechanical systems (Table 1), he discovered that there was a "missing mechanical circuit element" [36].

Table 1: Electrical-mechanical analogy

Electrical Quantity	Mechanical Analog
Voltage, e	Velocity, v
Current, i	Force, F
Resistance, R	Lubricity (Inverse of Friction), $1/B$
Capacitance, C	Mass, M
Inductance, L	Compliance (Inverse of Stiffness), $1/K$

The correspondence between spring and inductor and between damper and resistor is perfect (Table 2). However, this is not the case for a capacitor. A mass can be considered analogous, but the capacitor has two terminals where the mass has only one: its center of mass. Therefore, a mass is actually the equivalent of a grounded capacitor, and the missing element is the mechanical equivalent of a capacitor, with two terminals.

Table 2: Mechanical-electrical analogous equations [36]

Mechanical	Electrical
 $Y(s) = \frac{k}{s}$ $\frac{dF}{dt} = k(v_2 - v_1)$ spring	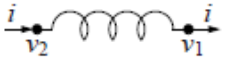 $Y(s) = \frac{1}{Ls}$ $\frac{di}{dt} = \frac{1}{L}(v_2 - v_1)$ inductor
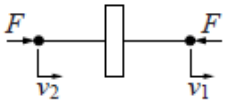 $Y(s) = bs$ $F = b\frac{d(v_2 - v_1)}{dt}$ inerter	 $Y(s) = Cs$ $i = C\frac{d(v_2 - v_1)}{dt}$ capacitor
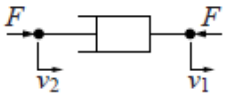 $Y(s) = c$ $F = c(v_2 - v_1)$ damper	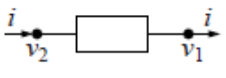 $Y(s) = \frac{1}{R}$ $i = \frac{1}{R}(v_2 - v_1)$ resistor

The concept of a system with two connectors in which the force is proportional to the relative acceleration has been theoretically proposed by Smith. The actual devices and patent [37] soon follow.

Inverters are actually very simple mechanical systems. The brilliant idea is converting the relative linear motion into a rotation and, through a gear set or not, spin a flywheel. Thus, with a fairly light system (around 1kg), it is possible to have an inertance of several hundreds of kilograms.

There are two main designs. The first uses a rack and pinion system (Figure 2.12).

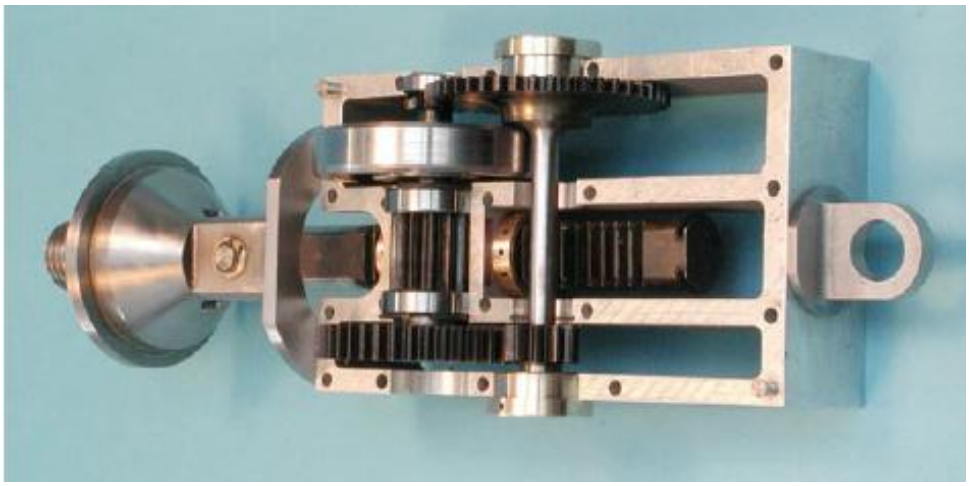

Figure 2.12: Rack and pinion inerter

The second design uses a ball screw and is the object of a second patent [38].

Figure 2.13: Ball screw inerter

A hydraulic inerter has also been developed [39]. A cylinder pushes some fluid through an external circuit, which incorporates a turbine connected to a flywheel.

Probably the most publicized application of the inerters is their implementation in the suspension of the McLaren MP4-20 Formula One car. In 2005, the F1 team started to use J-dampers (a decoy name to lure the competition) which, tuned correctly, improve handling and tire grip. Ride comfort (although not the primary concern in Formula One) is also an aspect that can be improved with inerters and is useful for everyday cars.

It would be possible to replace the flywheel of an inerter with the rotor of a generator, transforming the system into an energy harvesting device. This idea is similar to a design that will be presented later.

2.5 Maxwell's Equations

The complete solution for most engineering problems involving electric machines requires the use of Maxwell's equations. James Clerk Maxwell (1831 – 1879), after whom these equations are named, was a Scottish physicist and mathematician who is regarded as the founder of modern electromagnetic theory [40]. He developed laws unifying all previously known results, both experimental and theoretical, on electricity and magnetism. His results were first presented in the four-part paper *On the Physical Lines of Force* published in 1861 and 1862. Referred to as "the most significant event of the 19th century" [41], this publication is regarded as being of the same magnitude as Newton's *Principia Mathematica* and Einstein's *Annus Mirabilis* papers. Voted third-greatest physicist of all time (in a poll by Physics World) [42], Maxwell is described by Albert Einstein as "the most profound and the most fruitful that physics has experienced since the time of Newton" [43]. However, while Maxwell is ranked alongside Isaac Newton and Albert Einstein, his work failed to capture the popular imagination in the same way [44].

Modern references to "Maxwell's equations" refer to the set of the four partial differential equations as restated by Heaviside, leading to what are sometimes called the Maxwell-Heaviside equations [45]. The time-varying expressions of those equations in the differential and integral form can be seen in Table 3.

Table 3: Maxwell's equations

Law	Differential form	Integral form
Gauss's law	$\nabla \mathbf{D} = \rho_V$	$\oiint_S \mathbf{D} \cdot dA = \mathbf{Q}_V$
Gauss's law for magnetic fields	$\nabla \cdot \mathbf{B} = 0$	$\oiint_S \mathbf{B} \cdot dA = 0$
Faraday's law	$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}$	$\oint_{\partial S} \mathbf{E} \cdot dl = -\frac{\partial \phi_{\mathbf{B},S}}{\partial t}$
Ampère's circuital law	$\nabla \times \mathbf{H} = \mathbf{J} + \frac{\partial \mathbf{D}}{\partial t}$	$\oint_{\partial S} \mathbf{H} \cdot dl = \mathbf{I}_s + \frac{\partial \phi_{\mathbf{D},S}}{\partial t}$

where ∇ is the divergence operator

$\mathbf{D}$ is the electric displacement field, also called electric field density

ρ_v is the free charge density (not including bound charge)

$\mathbf{Q}_v$ is the net free electric charge

$\mathbf{B}$ is the magnetic field, also called the magnetic flux/field density

$\nabla \times$ is the curl operator

$\mathbf{E}$ is the electric field, also called electric field intensity

$\phi_{B,S}$ is the magnetic flux through the surface *S*

$\mathbf{H}$ is the magnetizing field, also called the magnetic field (intensity)

$\mathbf{J}$ is the free current density

$\mathbf{I}_S$ is the net free electrical current

Let's recall that in linear material: $\mathbf{D} = \varepsilon \cdot \mathbf{E}$ and $\mathbf{H} = \mu \cdot \mathbf{B}$

2.5.1 Gauss's Law

The Maxwell-Gauss equation, also known as Gauss's law, is expressed in its integral form as:

$$\oiint_S \mathbf{D} \cdot dA = \mathbf{Q}_V \tag{2.3}$$

This law states that the flux $\Phi_{\mathbf{D},S}$ of the electric displacement field $\mathbf{D}$ through a closed surface *S*: $\Phi_{\mathbf{D},S} = \oiint_S \mathbf{D} \cdot dA$ is equal to the free charge $\mathbf{Q}_V$ contained inside the volume *V* enclosed by that surface.

Gauss's law can similarly express the relationship between the electric field $\mathbf{E}$ and the total (free and bound) charge $\mathbf{Q}$:

$$\oiint_S \mathbf{E} \cdot dA = \frac{\mathbf{Q}}{\varepsilon_0} \tag{2.4}$$

where ε_0 is the electric constant.

Gauss's law provides a convenient means to determine the electric fields $\mathbf{E}$ and $\mathbf{D}$ for symmetrical charge distributions [40].

Gauss's law is an alternative way of expressing Coulomb's law, which states that the force *F* between two point charges $\mathbf{Q}_1$ and $\mathbf{Q}_2$ is:

- along the line joining them,
- directly proportional to the product $\mathbf{Q}_1\mathbf{Q}_2$ of the charges,
- inversely proportional to the square of the distance R between them.

Mathematically, it is expressed as:

$$F = \frac{1}{4\pi\varepsilon_0}\frac{\mathbf{Q}_1\mathbf{Q}_2}{R^2} \tag{2.5}$$

The application of the divergence theorem to Coulomb's law results in Gauss's law [40].

2.5.2 Gauss's Law for Magnetism

Maxwell's second equation appears under different names: Gauss's law for magnetism, absence of free magnetic poles, or the Maxwell-Thomson equation. It states that the magnetic flux density $\mathbf{B}$ is conserved. This means that no net flux enters or exits a closed surface [46]. It can be expressed by the integral equation:

$$\oiint_S \mathbf{B} \cdot dA = 0 \tag{2.6}$$

This implies that, as one of the names suggests, free magnetic poles, also called magnetic monopoles, cannot exist. The sources of the magnetic field will undoubtedly be dipoles, like magnets with a North pole and a South pole. It is not possible to separate the two opposite poles: if one would cut a magnet into two halves crosswise, this would result in two new magnets, both with North and South poles. In no case, will it end up with a pole on one side and the opposite pole on the other piece.

Another implication is that the magnetic field lines always form closed loops (or extend to infinity in both directions). The resulting pattern of a magnet, for example, is fairly well-known. It can be obtained with magnetic software (Figure 2.14) or can be revealed by iron fillings laid along a magnet as they arrange themselves along some field lines.

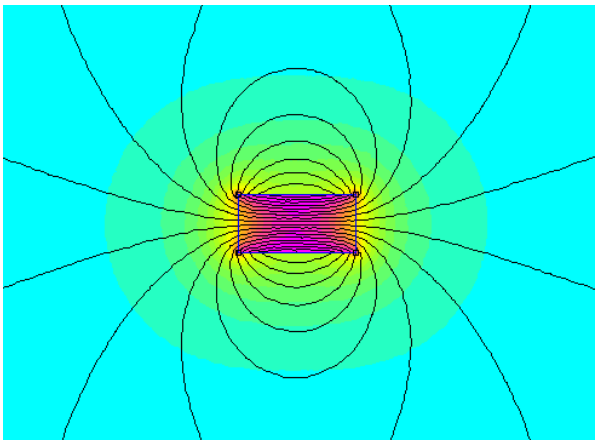

Figure 2.14: Magnetic field lines generated by a magnet surrounded by air

This property is important to remember when designing a system which involves magnetic fields, as the closed magnetic circuits will require particular attention for an efficient design.

2.5.3 Faraday's Law of Induction

After the experimental discovery by the Danish physicist Oersted that a steady electric current produces a magnetic field (which Ampère used to develop his law), the next step was logically to find out that reversely, magnetism can produce electricity. In 1831, Michael Faraday discovered that a time-varying magnetic field produces a voltage. Faraday's experiments showed that a static magnetic field generates no current flow, but a time-varying field creates an induced voltage, called *electromotive force* (emf), which causes a current in a closed circuit [40].

The Maxwell-Faraday equation, which is often called Faraday's law, can be expressed as:

$$\oint_{\partial S} \mathbf{E} \cdot dl = -\frac{\partial \phi_{\mathbf{B},S}}{\partial t} \tag{2.7}$$

That equation states that the line integral of the electric field intensity $\mathbf{E}$ along the closed contour ∂S is equal to the time derivative of the magnetic flux $\phi_{\mathbf{B},S}$ passing through that contour. That magnetic flux $\phi_{\mathbf{B},S}$ passing through ∂S equals the surface integral of the normal component of the magnetic field $\mathbf{B}$ over any non-closed surface S that has ∂S as boundary. Hence,

$$\phi_{\mathbf{B},S} = \iint_S \mathbf{B} \cdot \mathbf{da} \tag{2.8}$$

In electromagnetic machines with high electrical conductivity, like a generator for instance, the electric field $\mathbf{E}$ in the wire is extremely small and can be neglected [46]. The integral of the electric field is then equal to the induced voltage:

$$V_{emf} = -\frac{\partial \phi_{\mathbf{B},S}}{\partial t} \tag{2.9}$$

In the case of a coil of wire passed through by the exact same flux $\phi_{\mathbf{B},S}$, the electromotive force at the terminals of the coil is expressed as:

$$V_{emf} = -N\frac{\partial \phi_{\mathbf{B},S}}{\partial t} \tag{2.10}$$

where N is the number of (identical) turns in the coil.

The negative sign reflects what is known as Lenz's law. The current in the circuit that can arise from the electromotive force flows in such a way that the then-induced magnetic field will oppose the initial magnetic field.

Faraday's law gives an interpretation for the linear relationship (for a given flux) between the voltage and the speed often used to model electric permanent magnet motors and generators:

$$V = k_e \cdot \omega \tag{2.11}$$

2.5.4 Ampère's Circuital Law (with Maxwell's Correction)

From Oersted's experiments, Ampère develops a relationship between a magnetic field and current. In the integral form, it can be expressed as:

$$\oint_{\partial S} \mathbf{H} \cdot dl = \mathbf{I}_s \tag{2.12}$$

This equation is the magnetostatic equivalent of Gauss's law [40]. It states that the line integral of the tangential component of the magnetic field intensity $\mathbf{H}$ along the closed path ∂S is equal to the total current passing through that contour [46], i.e. crossing any non-closed surface S having ∂S as boundary:

$$\mathbf{I}_s = \iint_S \mathbf{J} \cdot \mathbf{da} \tag{2.13}$$

where $\mathbf{J}$ is the current density.

Although exact in magnetostatics, Ampère's law is actually inaccurate with time-varying conditions if Maxwell's correction is not included. The Maxwell-Ampère equation in the integral form can be written as:

$$\oint_{\partial S} \mathbf{H} \cdot dl = \mathbf{I}_s + \mathbf{J_D} \tag{2.14}$$

where the displacement current $\mathbf{J_D}$ is equal to $\mathbf{J_D} = \frac{\partial \mathbf{D}}{\partial t}$ where $\mathbf{D}$ is the electric displacement field (recall that $\mathbf{D} = \varepsilon \mathbf{E}$ in linear materials).

However, with electromechanical machines, the assumption can be made that considering the ranges of frequencies and sizes involved, the displacement current term can be neglected. This term will be required when dealing with electromagnetic radiations and time-varying electric fields in space [46]. The result of this approximation is the permanent or quasi-static form of the Maxwell-Ampère equation (2.12), which gives the particularly useful relationship between the magnetic field and the current from which it originates.

2.6 Lorentz Force – Laplace Force

The Lorentz force equation is often associated with Maxwell's equation as it makes the link between mechanics and electricity. According to the work of Lorentz, the force F exerted on a particle of charge Q in the presence of an electric field $\mathbf{E}$ and a magnetic field $\mathbf{B}$ can be expressed as:

$$F = Q(\mathbf{E} + \mathbf{u} \times \mathbf{B}) \tag{2.15}$$

where $\mathbf{u}$ is the speed of the particle [46].

This equation holds for one charged particle, but it can easily be modified to describe the force due to a magnetic field exerted on a piece of wire, in which a large number of charged particles (electrons) are in motion. The elemental force dF can be rewritten as:

$$\begin{aligned} dF &= dQ\mathbf{u} \times \mathbf{B} \\ \Rightarrow dF &= Id\mathbf{l} \times \mathbf{B} \end{aligned} \tag{2.16}$$

since $I\,d\mathbf{l} = dQ\mathbf{u}$ [40], where I is the current and $d\mathbf{l}$ is an infinitesimal piece of wire.

If the current is flowing in a circuit L, the force $\mathbf{F}$ on that circuit due to the magnetic field $\mathbf{B}$ is:

$$\mathbf{F} = I\int_L d\mathbf{l} \times \mathbf{B} \tag{2.17}$$

This shows that for a circuit in a given magnetic field, the force is proportional to the current. In a rotating device, this would lead to the linear relationship between torque and current.

However, the torque of a DC motor is not due to Lorentz force but to the repulsive force acting between the magnets on the rotor and those on the stator, whether they are permanent magnets or electro-magnets, depending on the type of motor. The electromagnetic torque can be expressed in terms of the flux per pole and the magnetomotive force. Assuming a constant magnetic flux, this leads to the classic linear relationship between the torque and the current used for DC permanent magnet machines:

$$T = k_t \cdot I \tag{2.18}$$

This page intentionally left blank.

Chapter 3

Linear Energy Harvesting Systems

3.1 Introduction

With an important background at the Center for Vehicle Systems and Safety (CVeSS) in the field of suspensions and dampers, the focus of this work from the start was on using the motion of the suspension as input for the onboard generator. From the simple observation that a suspension moves linearly, the idea came that a linear generator would be more appropriate. Using the exact same motion as the input removes the need for a complex mechanical system where efficiency would be lost. The concept is simple: an arrangement of magnets is fixed to one side of the suspension and a coil of wire is attached to the other side. When the train is in motion, the suspension moves, and so does the generator. From this starting idea, the hard part is to design an efficient system with the appropriate magnetic circuit, the proper coil, and all the correct parameters to be able to supply sufficient voltage (and power) to a battery.

For this application, a tuned vibration generator was considered but was quickly disregarded, primarily because the frequency encountered is not constant and can greatly vary depending on the conditions. However, the range could be narrowed down around particular frequencies, such as the natural frequency of the car body or that of any other bogie elements. The energy content is mostly confined to fairly narrow frequency bands. To that regard, a tuned vibrator, although not ideal, could work. However, an issue with a freight car is that the dominant frequency in the suspension is between 1 and 2Hz, far below the usual range of these devices. Another issue is that the target is to produce around 10 Watts of power which is well above what vibration harvesters are usually capable of. So, a system that is not as dependent on frequency and in which the magnet can be forced to move is deemed a better option.

The idea is to design a system that would resemble a coil-over damper, so the shape and size will be somewhat constrained by this choice. All the prototypes will be cylindrical.

3.2 Magnetic Analysis

To determine the most promising concept, numerous ideas have been evaluated using the magnetic field computed by finite-element analysis software FEMM.

3.2.1 FEMM

FEMM, or Finite Element Method Magnetics, is a free (distributed under the Aladdin free public license) software for electromagnetics. It is designed to solve magnetostatic, time-harmonic magnetic, and electrostatic problems in two-dimensional, either planar or axisymmetric, domains. The last release of the software also addresses steady-state heat flow and current flow problems. FEMM consists of three parts: the interface shell, the mesh generator, and the solver [47].

The interactive shell includes a pre-processor and a post-processor. The first has a CAD-like interface that allows the user to lay out the geometry, and to define the material properties and the boundary conditions. The model can be built by using the integrated drawing tools (points, lines, circle arcs) or by importing the geometry directly from CAD software in the Autocad DXF format. Therefore, a system designed in SolidWorks can easily be transferred to FEMM. The post-processor can display the field solutions in the form of contour and density plots. The user can obtain the values of the field at any point, simply by clicking, and can also create plots or compute integrals of various variables along user-defined contours.

The solver interprets the data files generated to describe the problem, and obtains the values of the field throughout the domain by solving Maxwell's equations.

The Lua language is integrated into the interactive shell to add scripting and batch processing to FEMM. Lua is a complete lightweight open-source programming language designed as a scripting language, which allows the files to be opened in a text editor, like notepad. All the FEMM commands have been added to the Lua language so every action in the graphic interface can be executed in a script. This allows the automation of repetitive tasks.

It is also possible to make FEMM communicate with numerical software like Matlab via ActiveX. FEMM acts as an automation server that other programs can connect to as clients.

3.2.2 Model

The various design ideas are then drawn in FEMM to evaluate the magnetic field and thus compute an expected output voltage. FEMM is a two-dimensional program. The axisymmetry mode is used here due to the nature of the system designed to be packaged like a shock absorber. In this mode, the y-axis is used as the axis of symmetry, and the section of the model to test is drawn on the right-hand side. Anything in the left-hand side will bring an error and will prevent the solution from being computed.

The many designs evaluated have been drawn directly in the graphic interface using FEMM drawing modes (Figure 3.1). Although basic, they yield the quickest way of creating the model. Using CAD software such as SolidWorks to create the model would be impractical because one would have to build the three-dimensional model in SolidWorks, generate the 2D DXF file, and then import it to FEMM. At that point, small features like small radius filets can pose a problem in generating the mesh. Consequently, the geometry has then to be (slightly) reworked. So, building the model in FEMM from the start may be the best way to proceed at this stage.

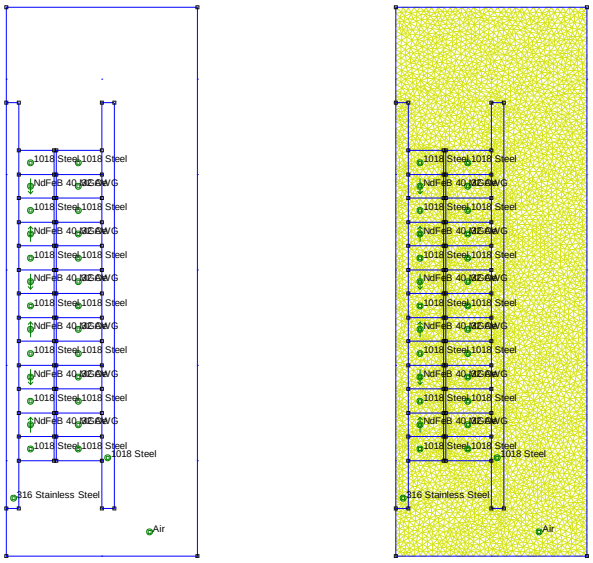

Figure 3.1: FEMM model (left) and meshed model (right)

3.2.3 Simulation

Once the model is built, simulations can be executed. FEMM solves Maxwell's equations for every node and computes all the parameters concerning the magnetic fields. The ultimate goal, beyond knowing the magnetic field, is really to estimate the output voltage of the system. As stated in Faraday's law (see Section 2.5.3), the voltage is equal to the opposite of the time rate of the flux. However, FEMM solves

magnetic problems that are fixed in space and in time. Without the possibility of considering time, FEMM is then incapable of computing the voltage. The problem then must be solved indirectly. The solution is to simulate the system with the rotor in various positions corresponding to a desired motion, and to record the flux for all those positions. It is then possible to link the flux to time so that the voltage with the derivative of the flux with respect to time can be computed.

The first step is to define the positions in which the model is simulated. Given the input frequency and amplitude, they are simply defined by the following equation:

$$z_i = A \cdot \sin\left(2\pi f \cdot t_i\right) + z_0 \tag{3.1}$$

where z_i is the position of the rotor at the ith step

A is the amplitude of the displacement

f is the frequency of the motion

z_0 is the starting position.

A script written in Lua (Appendix A) is used to automatically displace the rotor from one position to the next. In every position, the program commands FEMM to analyze and solve the model. It then computes the normal flux on predetermined paths inside the coils. The values of the flux are then presented in the console. The script can also ask to take screenshots of the output of FEMM, showing the magnetic field and field lines (Figure 3.2).

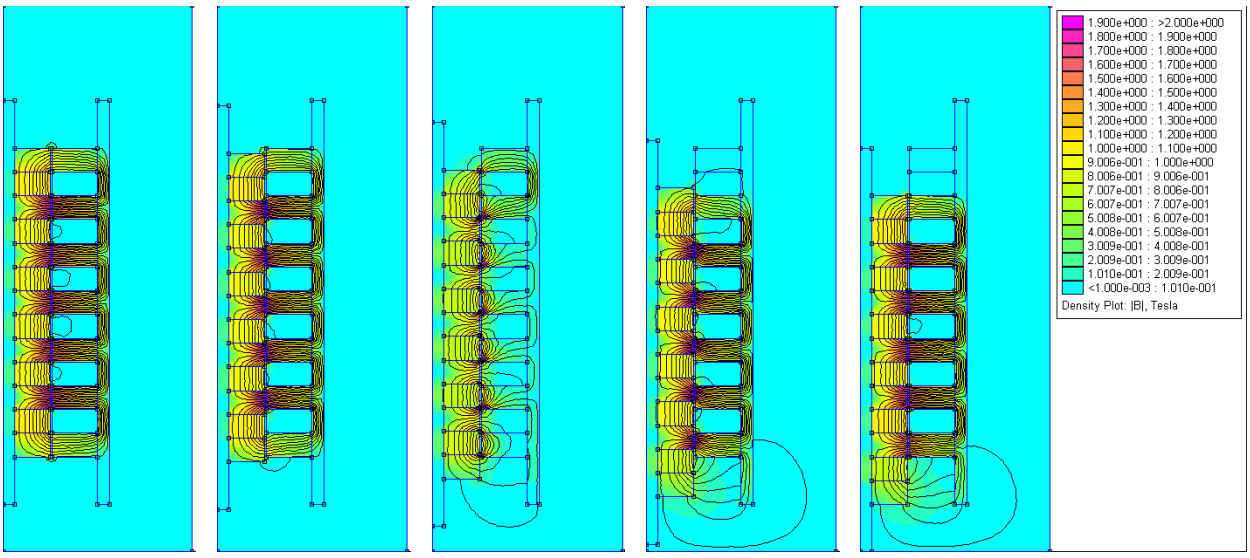

Step 1 Step 5 Step 10 Step 15 Step 20

Figure 3.2: Multiple steps of simulation using FEMM

The steps are usually equally spaced in time between the start and the end of the predefined number of oscillations.

Once all the steps have been executed, the console has all the values of the flux inside the coils, and the information can be extracted for post-processing and voltage estimation.

3.2.4 Post-processing

The raw data is exported to numerical computing software (e.g. Excel) to be analyzed. Time, position of the rotor, and values for the magnetic flux are known from the FEMM simulation. The voltage for each position is then estimated by differentiating the normal flux:

$$V_i = -N\frac{d\Phi_B}{dt} \approx -N\frac{\Phi_{i+1} - \Phi_{i-1}}{t_{i+1} - t_{i-1}} \tag{3.2}$$

where V_i is the voltage at the ith step

Φ_i is the magnetic flux inside the coil for the step i

t_i is the step time

N is the number of turns in the coil.

The differentiation is approximated by the slope between the previous and the next point. It is assumed that the magnetic flux is the same for every turn of the coil. With each loop being in series, the output voltage is then the sum of each loop voltage, which equals the product of the voltage from one loop by the number of turns constituting the coil. This number is determined by the size of the wire chosen, and by the space available between two magnetic poles. For example, in the model presented in Figure 3.1, there is enough room to fit a coil of around 250 turns of 24 AWG wire.

A sinusoidal displacement of 0.5-inch peak-to-peak has been simulated (Figure 3.3). The voltage generated across the coil is estimated to reach 1.5 Volts at peak.

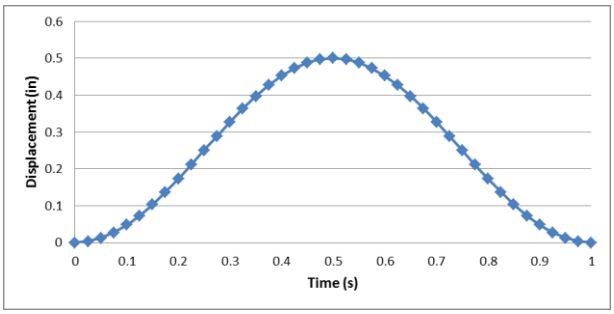

Displacement

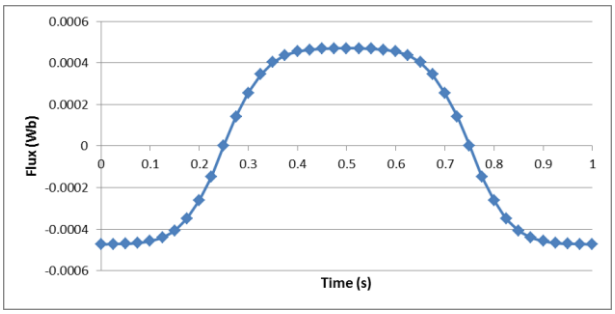

Flux (in one loop)

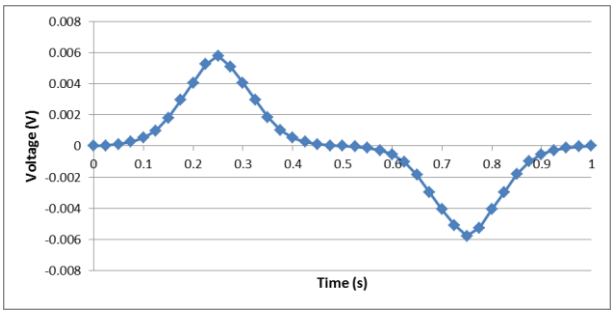

Voltage (from one loop)

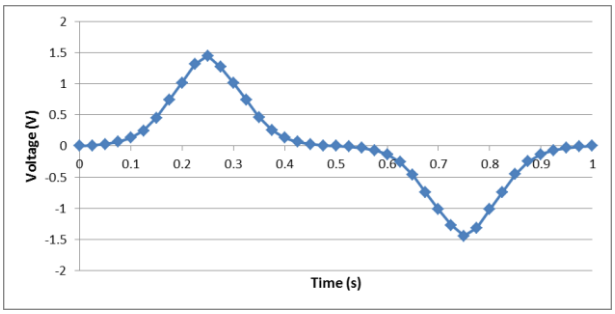

Voltage (from a coil – 250 turns)

Figure 3.3: Results extracted from the FEMM simulation

Over the entire cycle, the average voltage is 0.67 V_{RMS}. The current is driven by the load connected to the circuit. Thus, the current and therefore the power cannot be computed from this simulation, which effectively represents the open circuit voltage.

In Equation (3.1), which defines the displacement for each position, a term for the starting position is present. Simulations proved that this term is particularly important as it has a great impact on the output voltages, especially when small displacements are considered. The results presented in Figure 3.4 are from a simulation based on the exact same model used previously. The only difference is the starting point, which is a quarter of an inch higher. The motion and the rest of the simulation are identical to what has been performed before. It is easy to note the difference in the flux over time. However, if we correlate the flux to a position, we of course get back to the same values obtained before. At a given point, the flux will always be the same, independently of how the motion brings the rotor there. However, this yields important differences in the voltage. Indeed, if the rotor is moving fast in regions of rapidly varying flux, the voltage is fairly high. On the other hand, if the rotor is traveling slowly in those regions and fast in areas with almost constant magnetic flux, the output voltage is rather low. The latter case yields the plots shown in Figure 3.4. This can be explained by the following equation:

$$\frac{d\Phi_B}{dt} = \frac{d\Phi_B}{dz} \cdot \frac{dz}{dt} \tag{3.3}$$

It is clear that both the spatial variation of the flux and the time variation of the displacement need to coincide to obtain a large variation of the normal flux and thus of the output voltage.

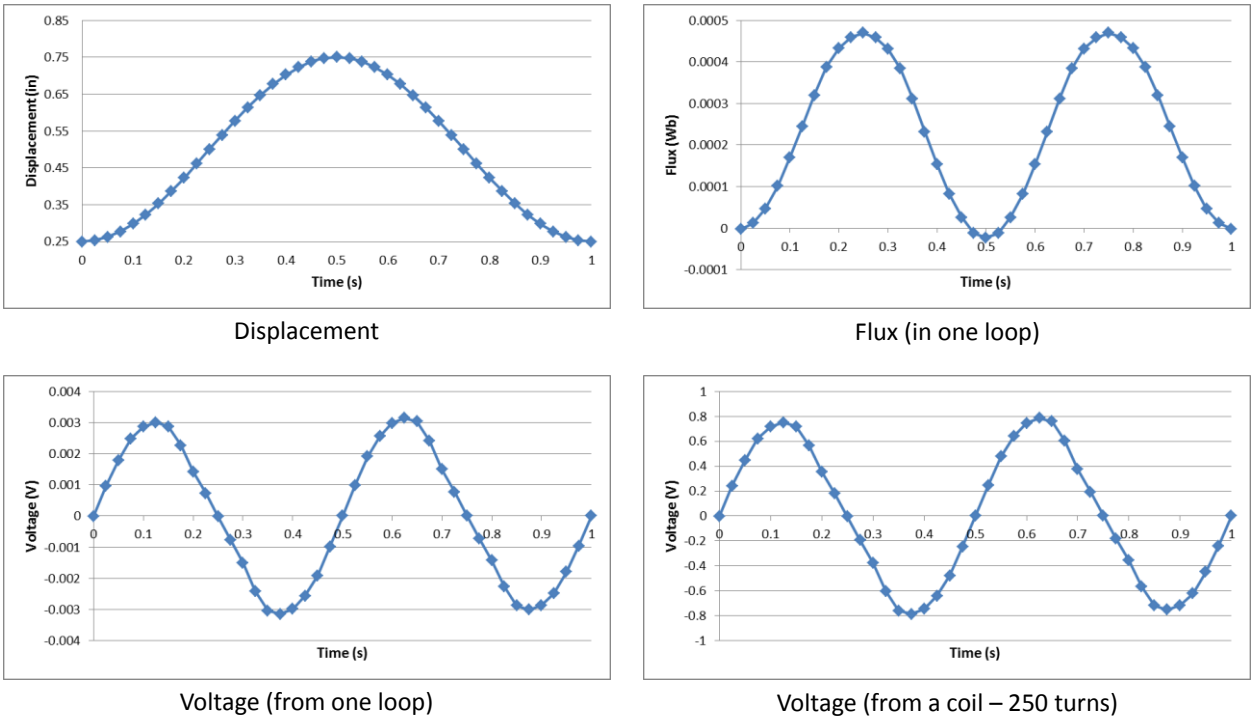

Figure 3.4: Results extracted from the FEMM simulation - Same model, different offset

The output voltage here only reaches 0.75 Volts, which is half of the first simulation. But instead of having two peaks, this displacement now offers four peaks per cycle. The average voltage generated is 0.53 V_{RMS}. This is not a decrease as significant as for the peak values, but it is still a 20% drop. This highlights the importance of the starting position of the rotor, and this is something to consider during simulations and later during the tests on the actual prototype.

3.2.5 Performance Criterion

The first goal of the simulations in FEMM was to design an efficient magnetic path. With the possibility of estimating the output voltage, the primary concern then shifted to produce a system able to generate voltages high enough to charge a battery from the relatively small amplitude and low frequency input. The voltages generated are, as mentioned before, linked to the changes in the magnetic flux. These voltages come from the motion of the magnetic rotor: the faster it moves, the larger the variations in the

flux and therefore the higher the voltage. Unfortunately, small amplitudes and low frequencies imply small speeds. Herein lies a real challenge for the design.

Power is another very important parameter of the generator. Unfortunately, FEMM does not allow for the possibility of computing power. The generator could be modeled by a voltage source, a resistor, and an inductor in series. The resistor represents the internal resistance of the generator simply due to the resistance of the wire used to make the windings. That resistance can easily be calculated based on the cross-section and the length of wire. The cross-section is an easy value to obtain. Wires come in standard sizes (gauges), with well-defined resistances per unit of length. The length of the wire is obtained by basic geometric computations. However, while the internal resistance is relatively easy to estimate, the inductance of the coil is not. In a first approach, one could neglect it, we will see later that considering the internal resistance does not allow a good match between the simulation and the tests.

The different configurations are then primarily compared on the RMS voltages they are expected to generate in open circuit.

3.2.6 Concluding Remarks

The entire process presented in this section is then repeated for all design ideas. An FEMM model is built for each concept and is then tested to evaluate the magnetic flux inside the coils. This yields an estimated value for the open-circuit voltage that the system would be able to generate. This process allows the ideas to be compared and promising concepts to be determined before entering the prototyping phase.

The main limitation of analysis using FEMM is the impossibility of accurately estimating the output power; however, the first step is to prove that useable voltages can be produced, even with the relatively small inputs offered in a train suspension.

3.3 Design Considerations

3.3.1 Materials

The choice of the materials is quite important as it will greatly impact the performance of the device. For applications such as motors, generators, or transformers, soft magnetic materials are used. The distinction

between soft and hard magnetic materials can be made by their hysteresis loops. As their coercive force indicates, a much larger reverse field is required for permanent magnet material than for soft magnetic material [48]. This is consistent with the use of hard materials as permanent magnets. On the other hand, soft materials usually have little magnetization left and need an electric current to be magnetized. In today's applications, iron-based alloys are the materials of choice. First, they possess the required magnetic properties: low loss, low coercive force, and high permeability [49]. In addition to magnetic properties, iron-alloys have good mechanical properties and are ductile enough to be formed or machined to the required shape. In numerous applications, the components are laminated to limit the formation of Eddy currents and lowered associated losses. Usually, the chosen material has to be ductile enough to be rolled in a thin strip of which the lamination is made [48]. Today, silicon steels are the soft magnetic materials used the most [49]. Increasing the silicon content raises the resistivity, which in turn lowers the core losses. Low carbon steels contain up to 0.25% of silicon, whereas silicon steels have over 1%. The most common content is about 3.25%. Steel containing about 6.5% silicon are also produced. They have good magnetic properties but their brittleness requires special care in manufacturing. Low carbon steels are also widely used for soft magnetic materials, especially in consumer-oriented products (small motors, generators, pole pieces, laminations for relays). Less expensive and easier to machine, they are a good alternative to silicon steels in applications in which the operating conditions are not extreme and the performance is not highly critical [48].

For its large availability and better machinability, low carbon steel is chosen as the soft magnetic material for the generator. The very common 1018 steel, with its 0.18% carbon content, is the material of choice for the application.

3.3.2 Air Gap

The air gap is the space between the pole and the armature. The gap is, by far, the most reluctant part of the magnetic path. The reluctance of air is much higher than that of steel and the entire reluctance of the magnetic path can be approximated by that of the air gap. Its size is then crucial, and from a magnetic standpoint, smaller is better. Increasing the air gap would lead to lower efficiency, and in a motor, more current to obtain the same torque. From a mechanical standpoint, a gap has to be present so the rotor can move freely inside the stator. Contact between those two elements can lead to decreased efficiency, premature wear, and even failure.

3.3.3 Permanent Magnet Materials

For low power (fraction of a horsepower) applications, permanent magnet machines are extensively used. Their main advantage is that they don't need an external excitation to create the magnetic field, which also makes them simpler in design compared to externally excited machines. They are also generally smaller, as the space for the magnets is less than for field windings and is often less expensive. On the other hand, they are constrained by the capabilities of the magnets, which are somewhat limited in the magnetic field they can produce. There is also a risk of demagnetization of the magnets either due to excessive currents in the windings or due to overheating [46]. However, the introduction of high-strength rare-earth magnets, like neodymium magnets, greatly reduced these restrictions.

Up to the 1970s, permanent magnets were made from Alnico or ceramic (also known as ferrite). Alnico 5, an iron-nickel-aluminum-cobalt alloy, has a relatively large residual flux but fairly low coercivity (Figure 3.5). Alnico 8 has a lower remanence, but its higher coercivity makes it less prone to demagnetization. Mechanically, they have the disadvantage of being brittle. Ceramic magnet materials (also called ferrite magnets) have significantly lower residual flux densities than Alnico materials but at the same time, have noticeably larger coercivity, improving resistance to demagnetization. The discovery of rare-earth magnet materials marked a turning point in permanent magnet technology, vastly expanding the field of possibilities. Made from a neodymium-iron-boron alloy, neodymium magnets have high residual flux densities, comparable to those of Alnico materials, but they also feature much higher coercivities: about 20 times higher than Alnico 5, and thus have a much larger maximum energy product [46].

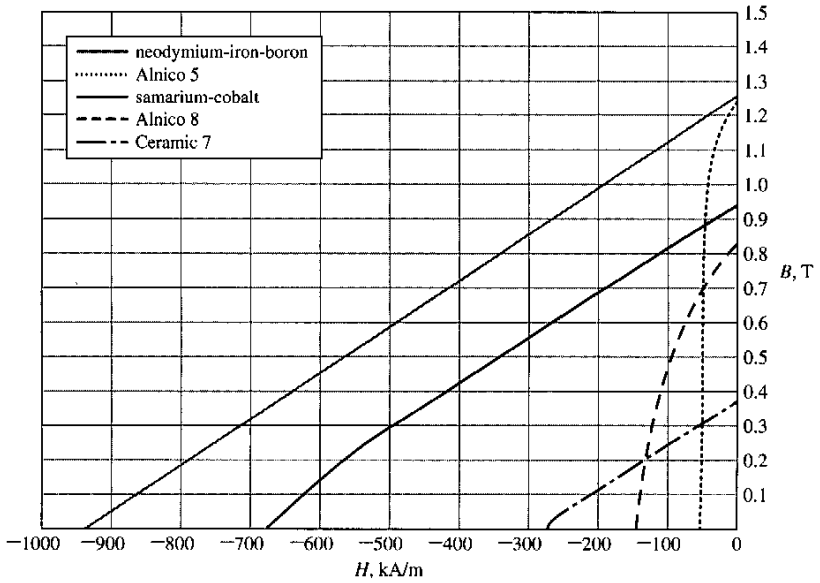

Figure 3.5: Magnetization curves of common permanent magnet materials [46]

Neodymium magnets have become widely available, in all size and shapes, and at affordable prices. They are the clear choice for the permanent magnets of the generator.

3.3.4 Permanent Magnets: Radial vs. Axial

As mentioned before, a voltage comes from the variations of the magnetic flux. So, the first element needed for the generator is a powerful magnetic core. The ability to create a high magnetic field is, with good relative velocities, the condition for large variations of the magnetic flux. Shaped like an elongated piston, the rotor is a stack of neodymium magnets and low carbon steel spacers. Strong magnets and highly permeable magnetic material is the right combination of material, but the configuration and magnetic path are also important factors in the design. The flux lines have to circle the coils in order to obtain the highest normal flux possible. Therefore, the flux out of the rotor must be radial.

Two options are then considered. The first option is to have radially-magnetized magnets, with a magnetic pole in the inside of the ring and the other pole on the outside face. The flux lines exit the magnets directly in the direction of the stator arms (Figure 3.6 - left). The second choice is to use ring magnets with poles on the flat faces. The flux travels axially into the steel spacer, where it is redirected outward (Figure 3.6 – right). Both configurations are tested in the finite-element magnetic software FEMM for comparison.

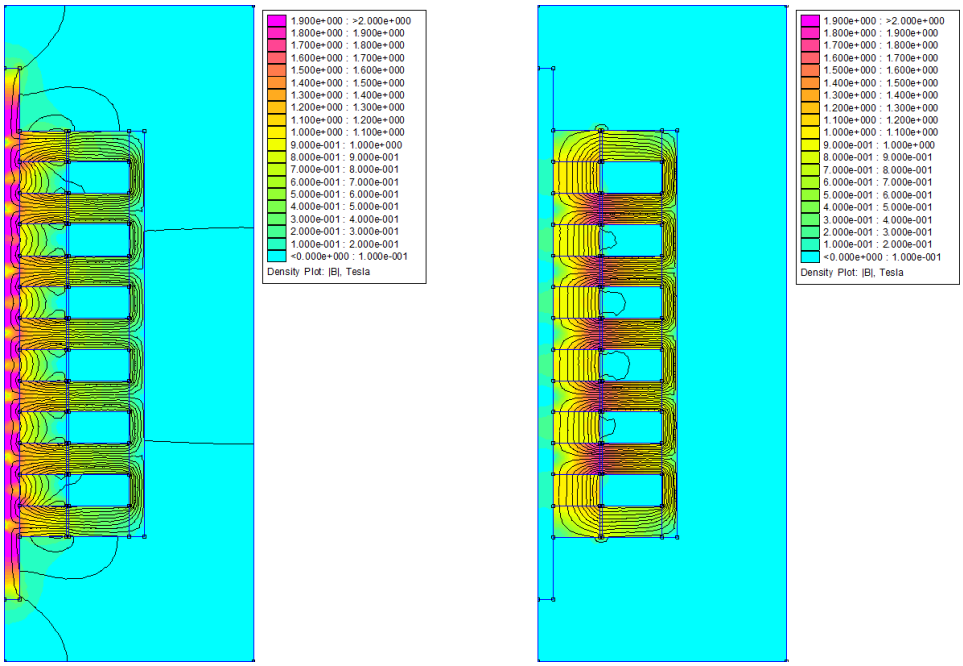

Figure 3.6: Side-by-side comparison of stators with radial magnets (left) and axial magnets (right)

The simulation result clearly shows the problem with using radially-magnetized magnets. The magnetic circuit passes through the inside rod where the cross-section area is relatively small. This bottleneck leads to saturation (purple areas on the figure) and reduces the efficiency of the system. A simple solution is to increase the cross-section of the magnetic path by using a rod with a bigger diameter. But, assuming the other dimensions are fixed, this would lead to smaller and therefore weaker magnets. A balance between rod diameter and magnet size needs to be determined. A series of magnetic simulations is performed to determine an optimal diameter for the rod. The system is tested with all parameters identical, but with a rod radius which varies from 0.025 inches to 0.475 inches (Figure 3.7).

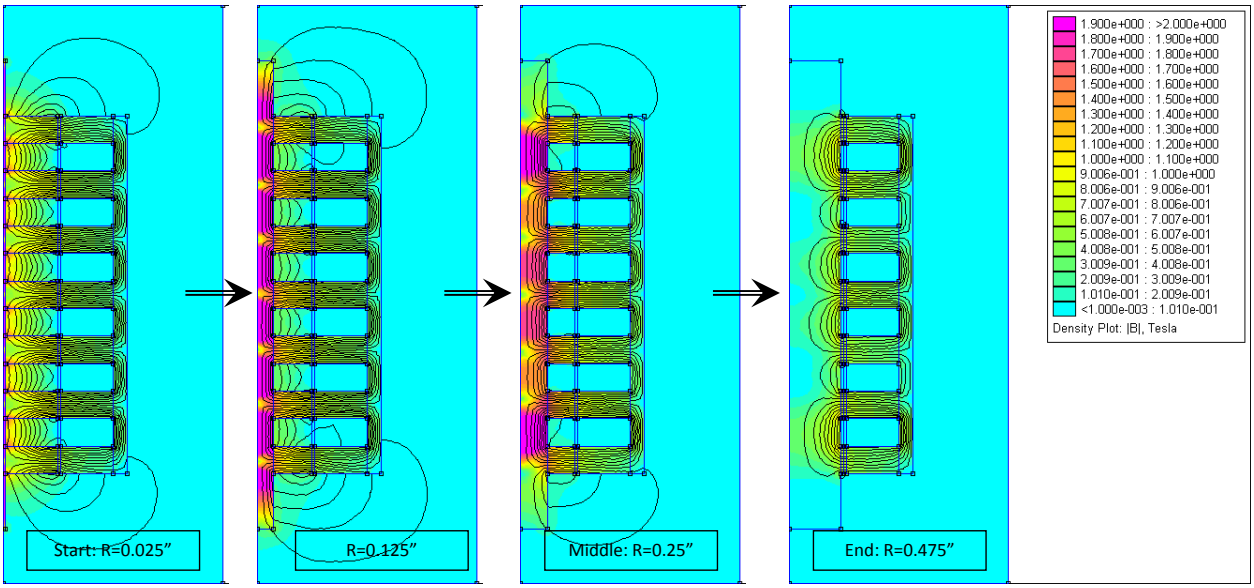

Figure 3.7: Simulation results for different rod radii with radial magnets

It can be noted that with a rod radius of 0.025 or 0.125 inches, the inner part of the steel spacers is crossed by some flux lines. That shows that this is part of the magnetic circuit and helps relieve the saturation. Ideally, with lower magnetic field, the flux lines should leave the magnet from its poles (like in the simulation with the rod of 0.25-inch diameter) and circle around the spacers, which could then be made out of non-magnetic material.

The normal magnetic flux inside the central coil is recorded for each simulation (Figure 3.8). It shows that initially, even if the magnets decrease in size, the flux inside the coil increases.

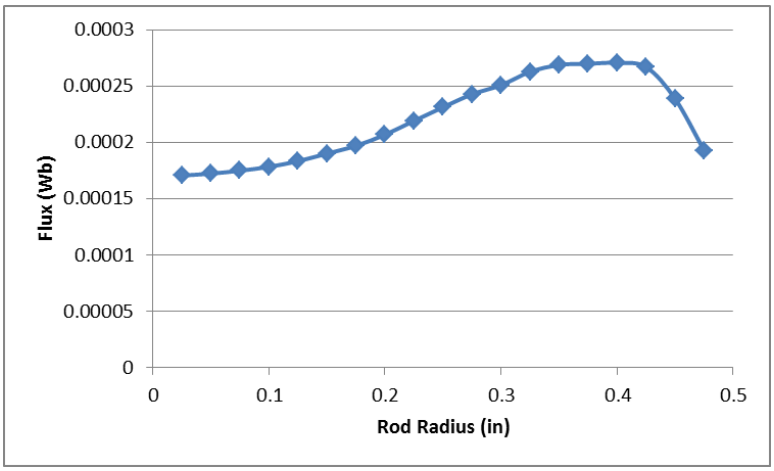

Figure 3.8: Normal flux inside a coil vs. rod radius for a system with radial magnets

This proves that the saturation in the rod limits the potential of the magnets. The flux reaches a local maximum before abruptly decreasing at the end. The final drop is due to the size of the magnet which becomes too small to create a strong flux and therefore take full advantage of the large diameter rod.

From this series of simulations, the optimal configuration is an inside rod with a diameter of about 0.75inch, for 1-inch outside diameter magnets.

With axially-magnetized magnets, the inside rod is not part of the magnetic circuit. It is even made out of non-magnetic material (such as stainless steel, brass, or aluminum) to prevent the flux from circling inside instead of around the coil, contributing to a higher normal flux and ultimately higher voltage. In that design, the small cross-section area on the magnetic path is then eliminated, and the areas of high saturation are no longer present. The circuit should be far more efficient. From a magnetic standpoint, the rod really has no use. It is only present to physically keep the rotor together when traction forces are applied. Although there is no optimum to determine, unlike the system with radial magnets, a similar series of simulations (Figure 3.9) with the rod diameter varying from 0.025 to 0.475 inches is performed to quantify the impact of the rod size, and to be able to compare the performance between the two types of magnets to ultimately make a decision on which one to use.

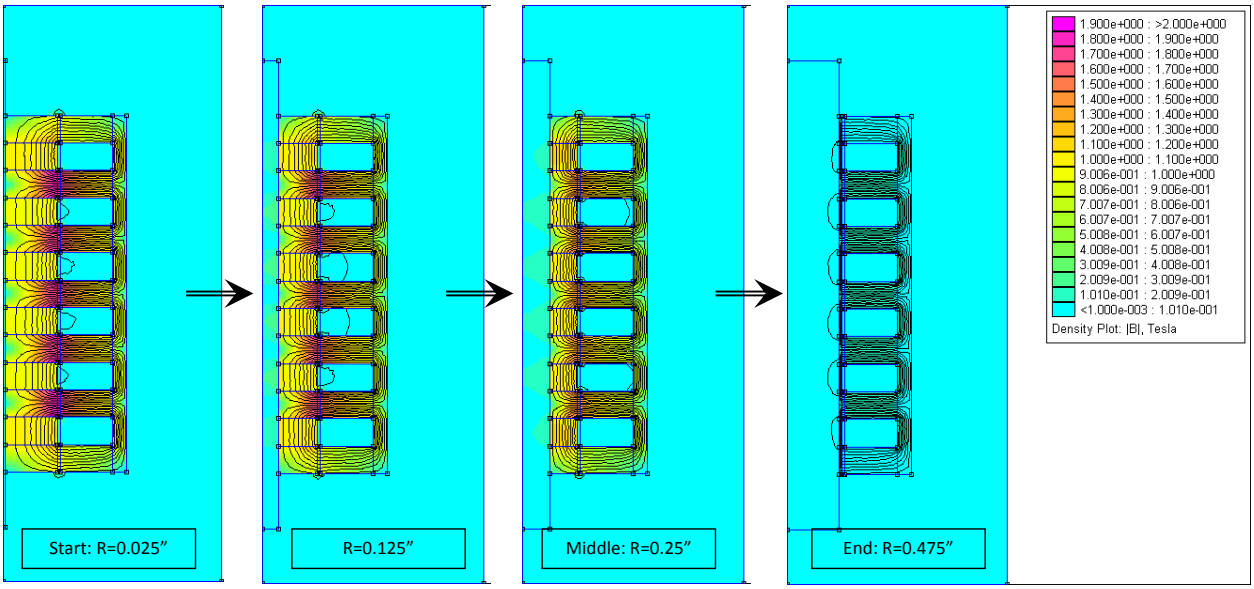

Figure 3.9: Simulation results for different rod radii with axial magnets

As can be expected, the normal flux inside the coil decreases as the rod radius increases (Figure 3.10). Unlike the previous simulations, the larger inside rod brings no magnetic advantage, and therefore the flux is monotonously decreasing. Furthermore, the quadratic shape of the curve is consistent with the proportional relationship between the magnetic field of a magnet and its volume. This definitively proves that for the best performance, the rod needs to be as small as possible.

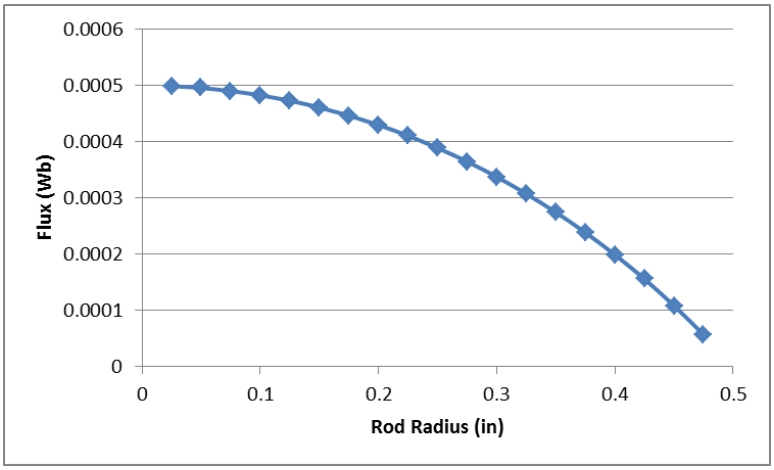

Figure 3.10: Normal flux inside a coil vs. rod radius for a system with axial magnets

With identical conditions (i.e. same stator and same dimensions) but different directions of magnetization for the magnets, the normal flux inside the coil in the axial magnet design is almost twice that of the radial magnet. This means that using the right magnets should yield roughly twice the output voltage. Axially-magnetized magnets are then naturally chosen to develop the generator.

3.3.5 Magnet/Spacer Dimension

The previous simulations show a high magnetic field, close to saturation, in the small regions around the air gaps. This is the new bottleneck, the smallest cross-sectional area, on the magnetic circuit. This region is an area to pay attention to in the design.

To quantitatively determine the impact of the width of the air gap, different configurations are tested. The size of the magnets and the coils remains unchanged with a thickness of 0.25 inch. Only the steel spacers between them vary: from very thin with only 0.025 inch, to fairly large at 0.625 inch (Figure 3.11).

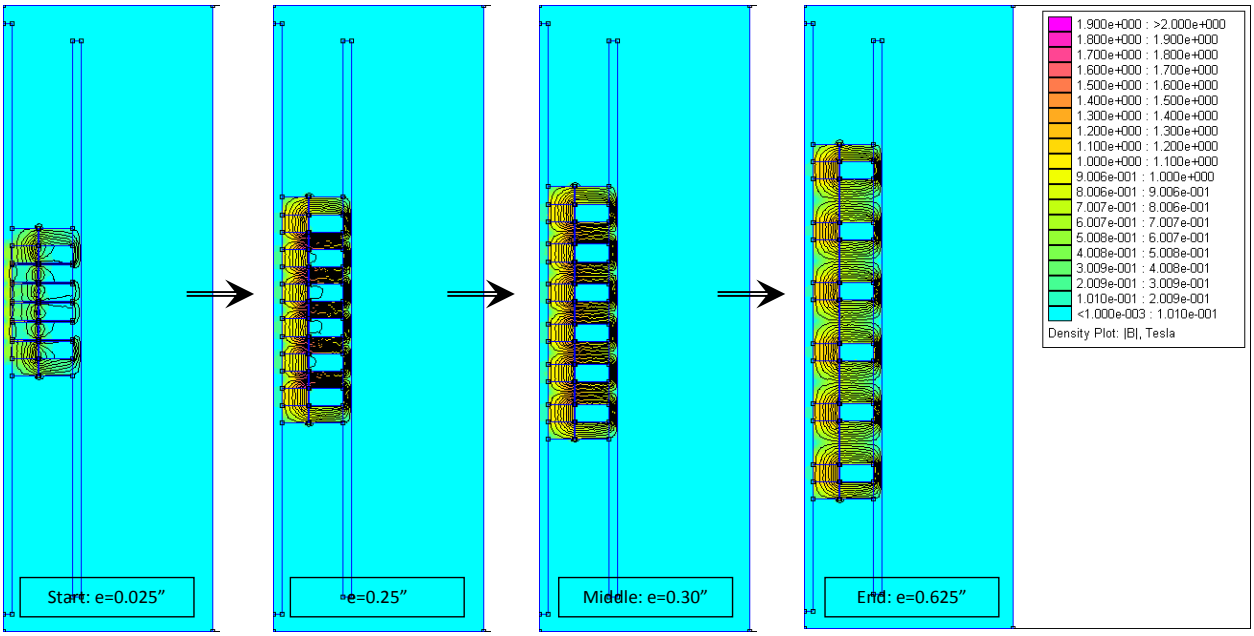

Figure 3.11: Simulation results for various spacer thicknesses

As can be expected, the normal flux inside the coils increases with the thickness of the spacers (Figure 3.12). The cross-section area of the magnetic path is less saturated and allows for a higher flux. We can distinguish two different portions in the curve: first, a rapid increase up to about 0.22 inch, followed by a slower rise. Although the larger spacers yield to higher normal flux, it will also lead to smaller variations

for the same displacements as the magnets are further apart. This is a particularly important aspect as the ultimate goal is not to have the highest flux possible, but the highest variations to yield to large output voltages.

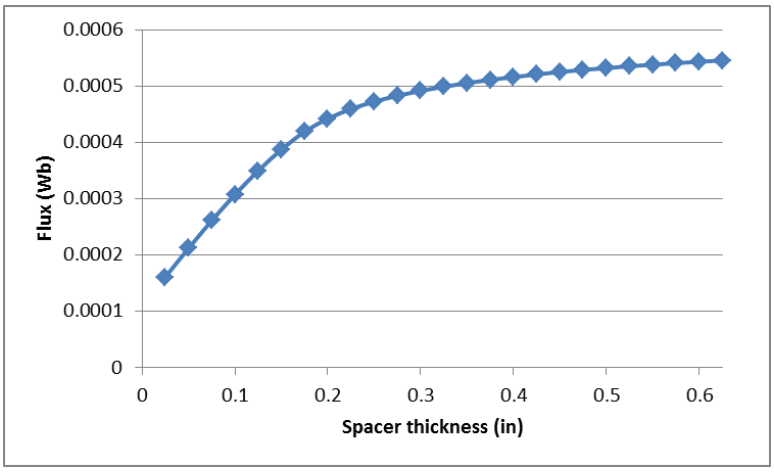

Figure 3.12: Normal flux inside a coil for various spacer thicknesses

In the same type of idea, the impact of the thickness of spacers only on the rotor can be evaluated. This comes from the observation that closer to the axis of the system, the cross-section area of the rotor spacer is smaller than that of the stator, for a given thickness. In that series of simulations, the stator remains unchanged and only the spacers in the rotor have their thickness increased (Figure 3.13).

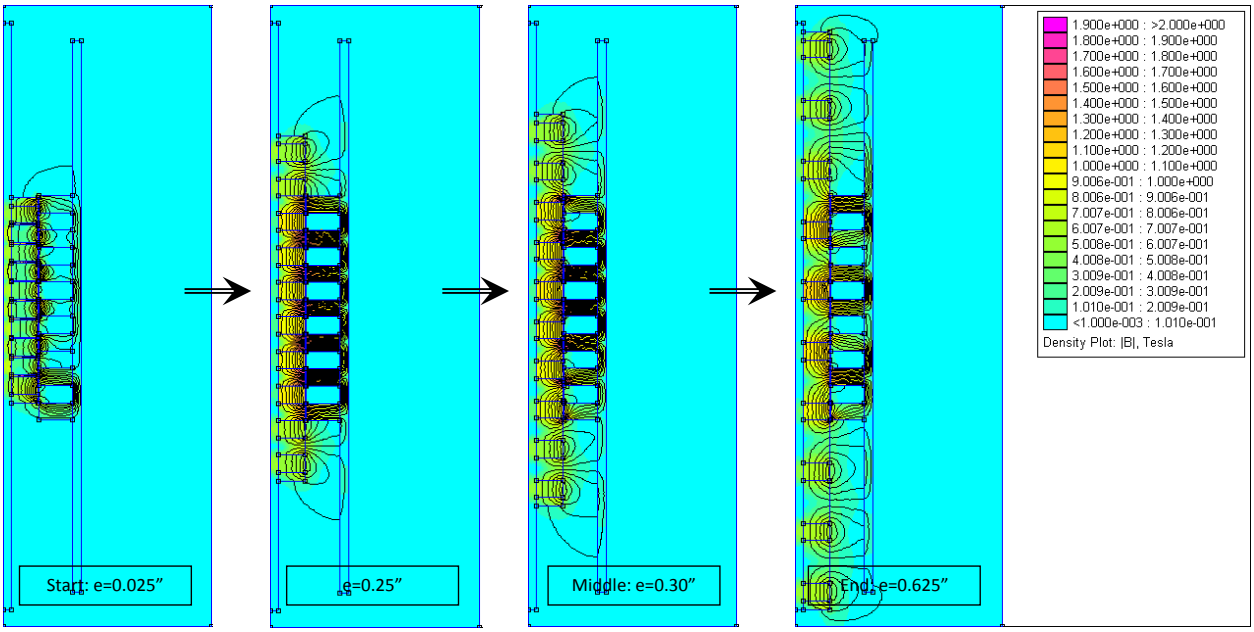

Figure 3.13: Simulation results for various rotor spacer thicknesses

Similarly to the previous results, the normal flux largely increases up to about 0.22 inch. But unlike previously, the flux then stagnates until the spacer is large enough to reach the size of two magnets (Figure 3.14). There is then no significant advantage to have the rotor spacers thicker than that of the stator.

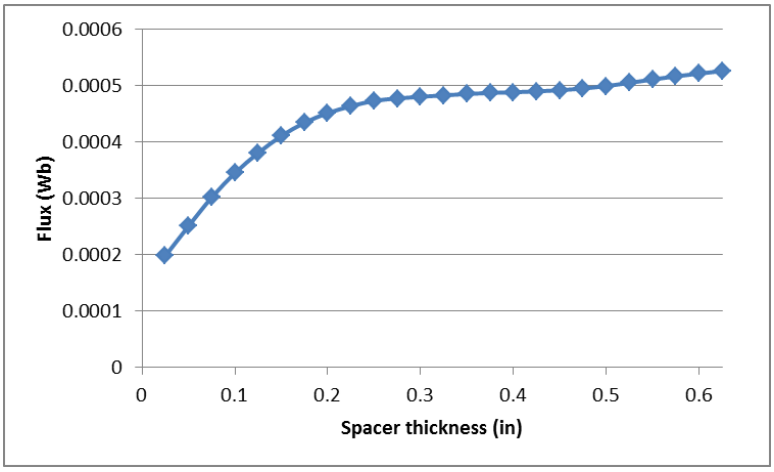

Figure 3.14: Normal flux inside a coil for various rotor spacer thicknesses

A spacer about the same size as the magnets seems to be a pretty good combination. In these discussions, magnets with an arbitrarily-chosen thickness of 0.25 inch are used. Different sizes would yield a similar conclusion. The difference in size only impacts the values of the normal flux, but since the comparison is relative, the conclusion remains the same.

The decision of the size of the magnets is based on geometric factors and on the expected input displacements. The diameter is governed by the space available. The system is meant to be inside the spring of the suspension, so the size of the generator is limited. The thickness of the magnets is also linked to the space issue to some extent but it is mainly a function of the input displacements. If the magnets are too big for the typical displacements across the suspension, the magnetic field will not vary by much. If they are too small, there will be more variations, but the amplitudes of the magnetic flux will be small and it can become difficult to fit a coil of a decent size. The magnets are then chosen to be 0.125 inch in thickness, so a displacement of only one-quarter of an inch would lead to a full cycle.

3.3.6 Numbers of Turns per Coil

As stated in Faraday's law, the output voltage for one loop is equal to the time rate of the magnetic flux. For a coil, it is equal to the sum of all the contributions. Assuming that all the turns of a winding are

subject to the same flux, the output voltage is proportional to the number of turns in the coil. Therefore, the larger the coil, the higher the voltage. However, this has its limitations. First, the space available usually limits the amount of wires that can be used. Also, the longer the wire, the higher the internal resistance and the higher the copper losses. This will be particularly important when current is drawn. With very large coil, the inductance can become an issue too. These aspects then have to be balanced with the desired output voltage.

3.3.7 Concluding Remarks

This section presented different aspects considered during the design of the generator. There are numerous parameters that play a role in the performance of the system, and analyzing them one by one allows us to make educated decisions. The elements that have been introduced, such as axially-magnetized neodymium magnets and low carbon steel parts, will then be used.

3.4 Designs

The design process went through numerous iterations, starting with the very basic idea used to present the process of how to estimate the voltage using FEMM in Chapter 3.2.

The rotor is an arrangement of disc-shaped permanent magnets with steel spacers in between. The same poles of the magnets are facing each other, such that the magnetic field is forced to exit the rotor radially through the steel spacers. A rod is placed through the whole assembly so traction efforts can be applied without pulling the magnet assembly apart. A shaft is attached to the through rod to be connected to one side of the suspension. The magnets were initially 0.25-inch thick, but to generate more variations for the same given displacement, the thickness is chosen to be one-eighth of an inch. The coils are then sized to match that dimension (Figure 3.15): the coils and the poles are both made 0.125inch wide.

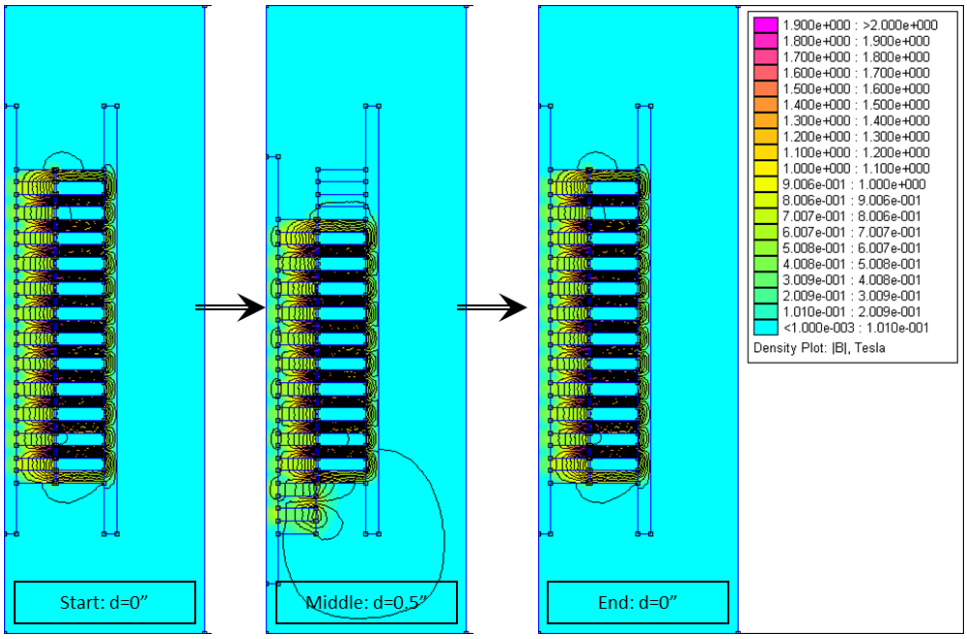

Figure 3.15: Simulation with 1/8" wide coils

With this design, the magnetic flux is fairly high and the variations are good even with displacements as small as one-quarter of an inch. The drawback is that the space for a coil is limited, and windings with only a limited number of loops can be implemented, without going into really thin wires. In this configuration, it is estimated that, with 200 turns, a coil would be capable of generating about $0.52V_{RMS}$ and $1V_{peak}$ for an input of ±0.25 inch at 1Hz.

Thus, the design moves to larger coils. The same rotor is used, with unchanged magnets. The gaps between two poles also remain the same, but their shape changes for a wider head. It can be understood as one every two coils grows larger whereas the neighboring coils just disappear (Figure 3.16).

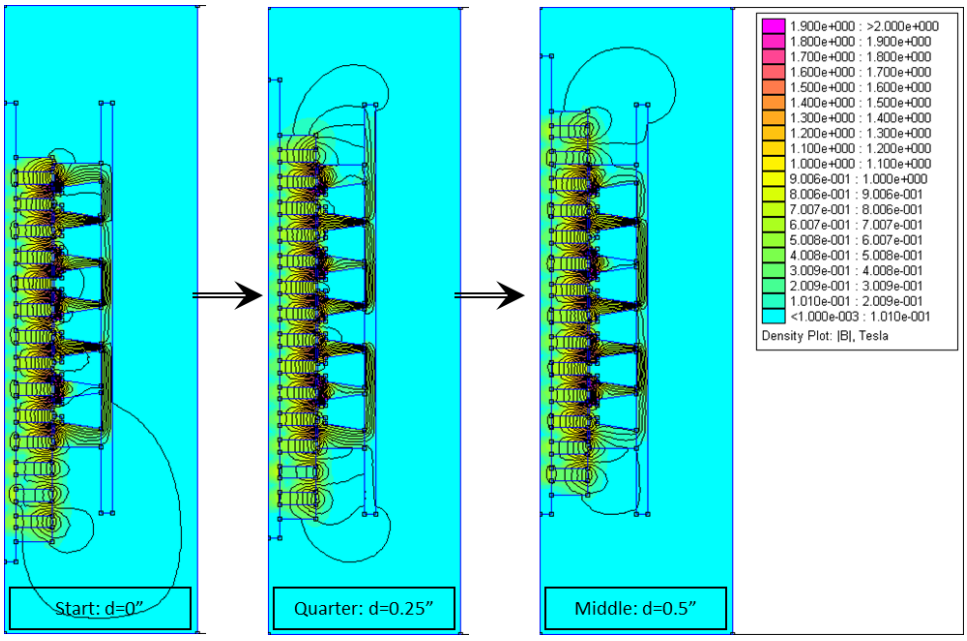

Figure 3.16: Design with wider coils

It can be noted in Figure 3.16 that the stator teeth are tapered, with thinner edges on the outside. The cross-section of the magnetic path is the important parameter and as it gets further away from the axis of symmetry, the circumference increases, leading to the possibility of a thinner part without compromising the magnetic path.

In this design, the coils are also phased, i.e. the spacing of the coils is different from that of the magnets. This yields the maximum of the magnetic flux reached at different moments for each coil, similar to what is found in a rotating motor. This allows the electromotive force to be spread over the displacement (Figure 3.17), and also reduces the cogging, i.e. the force that pushes the machine to rest in given positions, and gives a smoother motion.

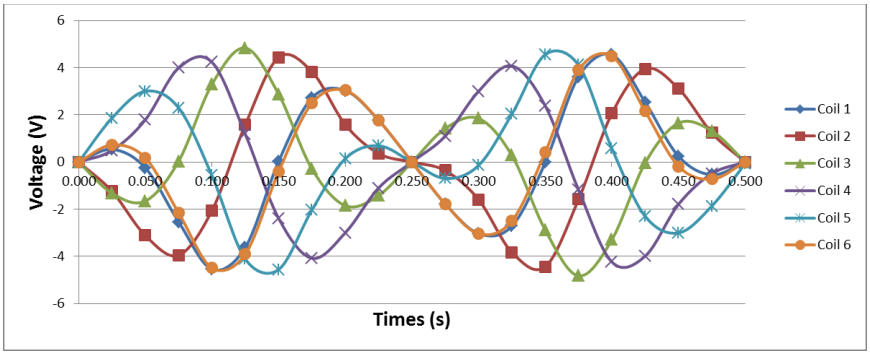

Figure 3.17: Output voltages with off-phase coils

With a displacement of ±0.25 inch at 1Hz, each coil (of 520 turns) is estimated to yield about 1.2-1.4V_{RMS} and around 2.5V_{max} at peak. At 2Hz, the outputs double, as expected. While still relatively low, this is an improvement over the previous design using smaller coils. Although the output would be smoother with those off-phase voltages, they cannot easily be combined to give rise to large output values. So, the spacing of the coils is changed to match that of the magnets on the rotor. In this configuration, with every other magnet used at a given time, the intermediate stator teeth don't appear to be on the magnetic path, which instead, encircles the entire coil assembly. This leads to the idea of a design using a single large coil (Figure 3.18). Openings in the inner core of the coil prevent the magnetic field from half the magnets to circle inside the coil, and oblige it to circle around it, leading to a high magnetic flux inside the coil.

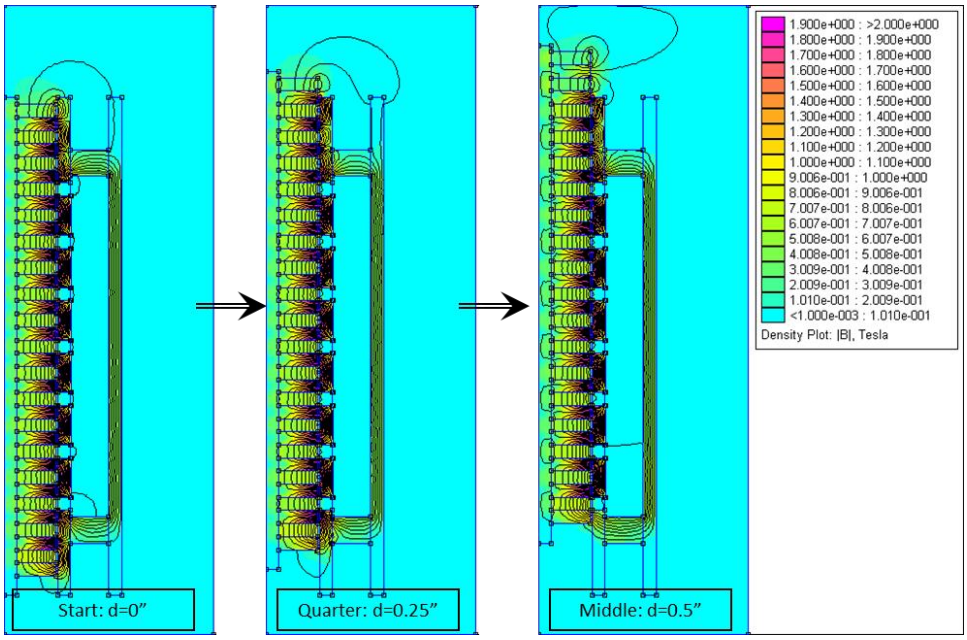

Figure 3.18: Design using a single large coil

With the teeth removed, there is a lot of space to place a large coil with many turns. This leads to an estimated voltage between 11 and 12 V_{RMS} for an input of ±0.25 inch at 1Hz.

A similar design including two large coils with a phase of 90° between them has also been investigated. The two phases would give a smoother output voltage but of lower value, thereby limiting the interest. An advantage of such a configuration would also have been the reduction of the cogging, similar to the multiple coil design discussed earlier.

3.5 Large Coil Design

3.5.1 Design/Idea

The voltage at the terminal of a coil of wire is the sum of the voltage generated by each individual turn, which is equal, according to Faraday's law (see Section 2.5.3), to the opposite of the variation of the flux through the surface surrounded by the wire loop. Assuming that every loop of the coil has the same flux, the output voltage is proportional to the number of turns. So, again, for voltage generation, more turns is better. However, this comes at the cost of increased internal resistance, which also needs to remain relatively low in order not to be an issue.

The challenge is then to place a coil with a large number of turns of decent-size wire in a limited space and in such a way that the magnetic path will circle around it. This last condition is important because if there are some loops of flux lines inside the coil, they will cancel each other out during the integral of the normal magnetic flux density over the surface bounded by the coil, and ultimately won't contribute to create a voltage. The idea is then to have a steel tube inside the coil that creates a magnetic path around the coil (Figure 3.19). An opening in that tube effectively forces the magnetic field the long way around.

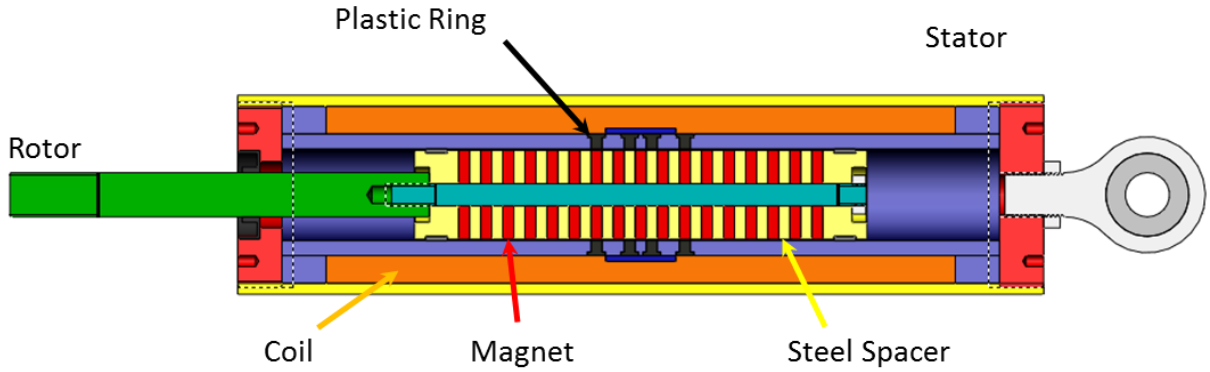

Figure 3.19: Large single-coil design

Unlike the design presented in the last section, this concept has only two openings to create the variable magnetic flux. Having openings all along the coil requires a longer rotor and it actually does not improve the performance as the parts of the magnetic path are already saturated. More excitation cannot increase the magnetic flux.

In Figure 3.19, two plastic rings can be noted in the middle section. Those do not contribute to improve the magnetic performance of the system, but to improve the cogging. The idea is that it gives a second

stable position to counterbalance the first one, and thus limit the force required to move the rotor. The force is not entirely canceled, as the attractive force of a magnet is not linear but proportional to the square of the distance.

3.5.2 Simulation(s)

As every other concept before, this design is evaluated through the same magnetic analysis using the finite-element software FEMM. The model is created and analyzed (Figure 3.20) in the different positions of a desired input: ±0.25in at 1Hz, for example.

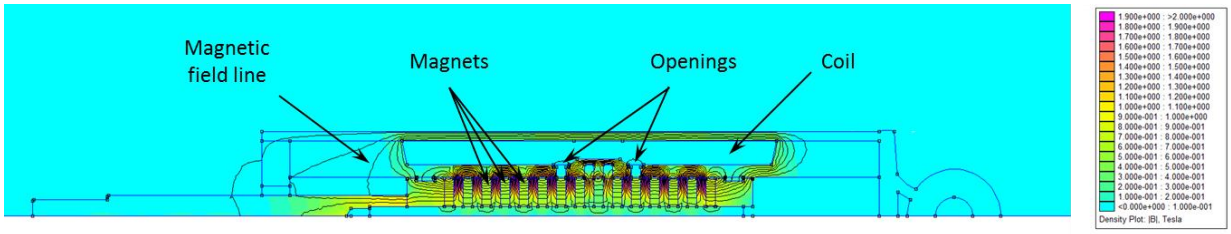

Figure 3.20: Magnetic finite-element analysis of the design

Every 50ms, the rotor is put into the predetermined position and the normal magnetic flux inside the coil is recorded at 11 equally spaced locations (Figure 3.21). The average flux is then used to compute the output voltage by differentiating with respect to time.

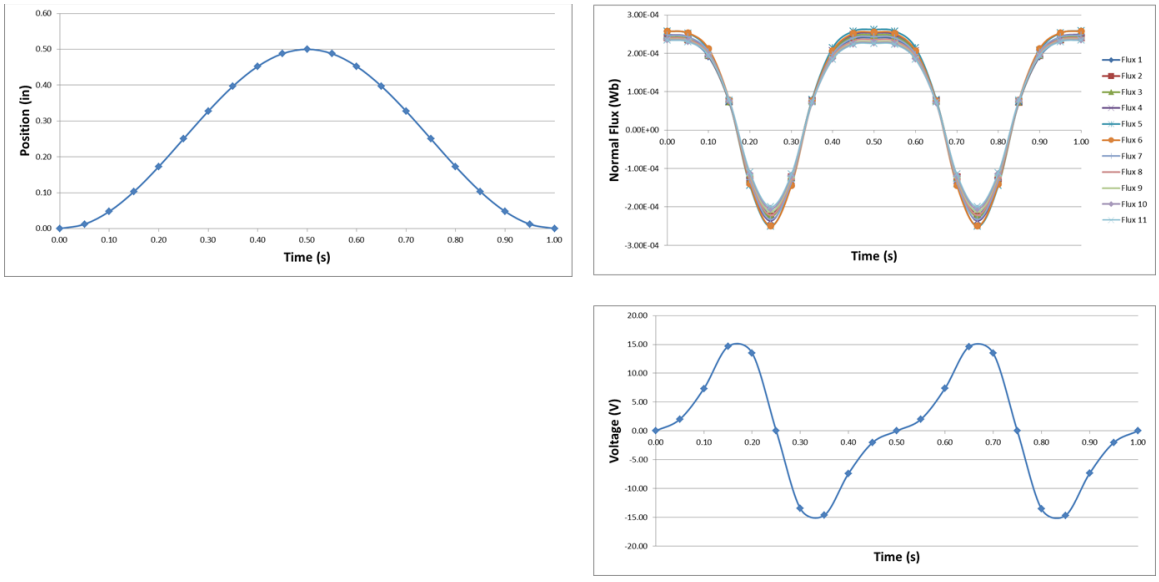

Figure 3.21: Simulation and post-processing results for the single coil design

The output voltage is then expected to reach over 15V at peak and to average a little under $10V_{RMS}$, for the small input of only one-quarter inch of amplitude at 1Hz. As mentioned earlier, the starting position is important in the output voltage, so other offsets are simulated (Figure 3.22). As expected, the profile shows some variations, with maxima at slightly different places. It is then anticipated that the generator will be able to produce an output voltage between 8.2 and $9.8V_{RMS}$.

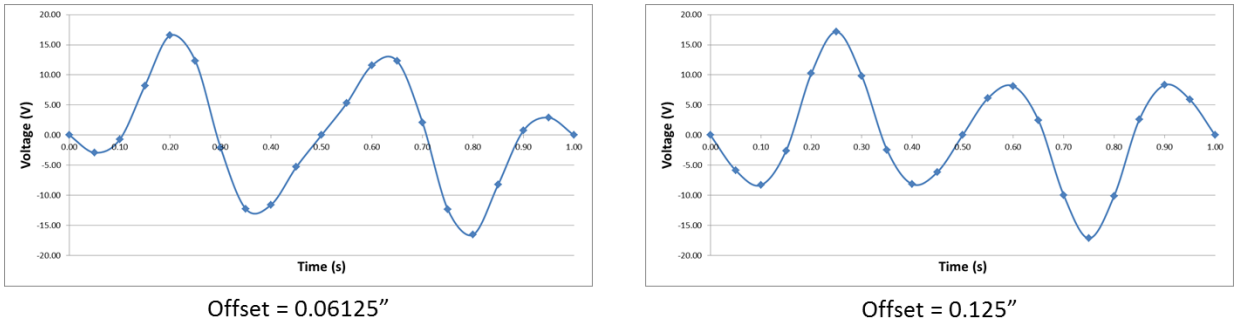

Figure 3.22: Voltage estimations for different starting positions

3.5.3 Prototype

With the details of the concept defined, a prototype was built. Most of the parts were machined by the department machine shop. After their reception, the first step was then to assemble the coil core, comprised of 10 parts (Figure 3.23).

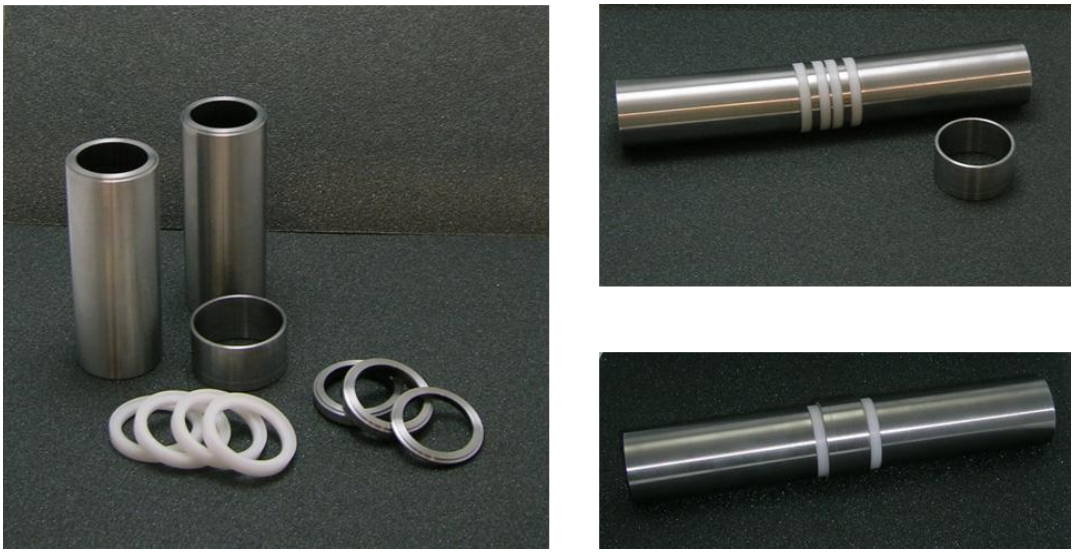

Figure 3.23: Coil core

With the inner core ready, the coil needed to be wound. The core was then placed on a round bar that can be clamped in the laboratory lathe (Figure 3.24). The winding was then slowly created by carefully placing the wire by hand. A counter kept track of the progress. Finished, the coil counts 4500 turns of 24 AWG wire. Electrical tape was placed around it to protect it and prevent the wire from becoming loose. It was then ready to be slid inside the outer tube.

Figure 3.24: Coil during winding and finished

The other part to be assembled was the magnetic rotor. It was comprised of a succession of 18 steel discs and 17 neodymium magnets. Special care had to be taken to make sure the magnet to be placed had the same pole facing the magnet already in place. The last part was the shaft to complete the rotor (Figure 3.25).

Figure 3.25: Magnetic rotor

The initial design was to have two wear bands, one at each end of the magnet stack. However, this presented two issues. The first is that it is nearly impossible to slide the rotor inside the coil in that way. After the first wear band is in the coil inner tube, the magnets cannot stay centered and stick to one side, preventing the rotor from being fully inserted. The rotor then has to be first inserted in a plastic tube which, held aligned with the coil, allows the rotor to be put in place. But then, the other issue is that the magnets are strong enough to bend the inner rod so they touch the coil core, creating a lot of friction. Perfectly centered, the magnets are at an equilibrium but a very instable one. The smallest deviation from the center will lead the rotor to move further and further away, up to contacting with the coil. The solution is to place a wear band all around the rotor, filling the air gap almost entirely. It is still easier to insert the rotor by using a plastic guide tube.

Once the rotor is in place, two screw-on caps seal the outer tube, completing the prototype (Figure 3.26). Two ball joints are present to be attached to the test rig.

Figure 3.26: Prototype fully assembled

3.5.4 Tests

The prototype was tested using the Roehrig EMA-2k shock dynamometer present at the CVeSS laboratory (Figure 3.27).

Figure 3.27: Prototype in test in the Roehrig shock dynamometer

It allows testing of any desired displacement input, and it records all the mechanical parameters: position, velocity, and forces. In addition to the mechanical aspect of the test, the output voltage of the system is recorded with a digital oscilloscope (HP 35665A). To estimate the power produced, different resistive loads are connected at the output of the system. The oscilloscope records the voltage across that resistor, and the power is computed using Ohm's law:

$$P = V \cdot I = \frac{V^2}{R}$$

The main disadvantage of this setup is that the two pieces of equipment recording the data are not synchronized and do not have the same sampling rate. This requires some extra steps during data processing.

The prototype was then tested with various inputs (different amplitudes and frequencies) and with different resistive loads. The recorded output voltages are summarized in Table 4: they range from $4V_{RMS}$ with the smallest motion and the lowest resistance, to close to 25VRMS for an input of ±0.75 inch at 1.5Hz and no load. The 10MΩ resistance is actually the internal resistance of the oscilloscope. This is then as close as possible to open circuit.

Table 4: Output RMS voltages

Input	180Ω 1in	180Ω 1.125in	2kΩ 1in	2kΩ 1.125in	10MΩ 1in	10MΩ 1.125in	Resistor Offset
±0.25in @ 1Hz	4.181	4.482	5.964	6.32	6.129	6.561	
±0.5in @ 1Hz	6.758	6.904	11.428	11.647	12.066	12.311	
±0.75in @ 1.5Hz	9.9	9.935	22.6	22.625	24.933	24.978	

It is important to note that the voltages go down as the resistive load values get smaller. This is due to the increased current drawn which induces higher losses (for example, ohmic losses). Table 5 shows that although the voltage decreases, the output powers can be higher due to an overall gain with the increased currents. In open circuit, no current flows, so the output power is evidently zero. With an 180Ω load, the average power range is from about 100mW to 550mW (Table 5).

Table 5: Average output powers

Input	180Ω 1in	180Ω 1.125in	2kΩ 1in	2kΩ 1.125in	10MΩ 1in	10MΩ 1.125in	Resistor Offset
±0.25in @ 1Hz	0.097	0.112	0.018	0.020	0.000	0.000	
±0.5in @ 1Hz	0.254	0.265	0.065	0.068	0.000	0.000	
±0.75in @ 1.5Hz	0.545	0.548	0.255	0.256	0.000	0.000	

For a given load, the average voltages are proportional to the average input velocities. So, with higher frequencies, the output values increase, as can be seen in Table 6. The average power is measured to reach over 3 Watts, with a 2kΩ resistor.

Table 6: Results for higher frequency tests – 2kΩ load

Input	Voltage (Vrms)	Power (W)
1in @ 3Hz	45.263	1.024
1in @ 5Hz	61.495	1.891
1in @ 8Hz	78.085	3.049

Beyond the average values, since the output is not DC, i.e. constant, it is important to look at the time signals (Figure 3.28). The output voltages are alternative and the combination of two frequencies: the envelope of the signal is due to the motion, whereas the higher frequency is generated by the magnet spacing.

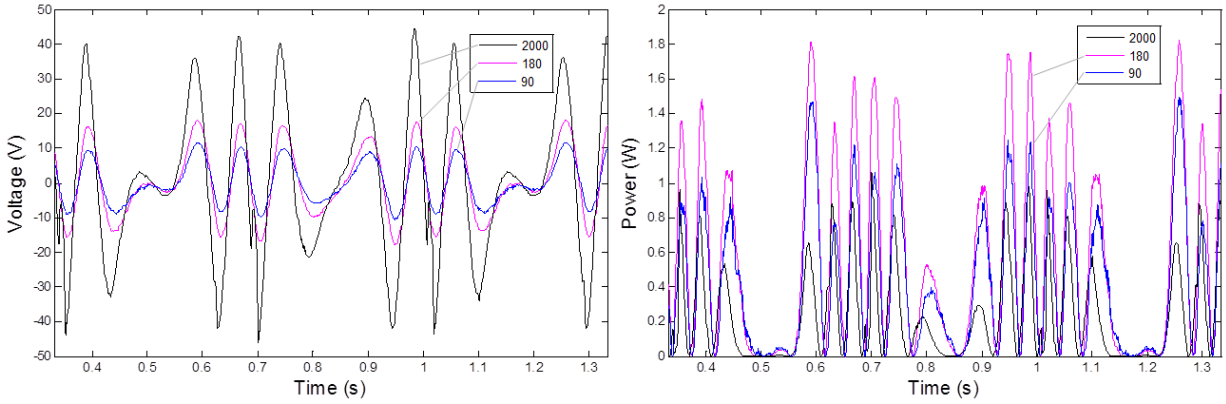

Figure 3.28: Voltage and power under different loading conditions - ±0.75" at 1.5Hz

In Figure 3.28, it can be noted that the power with an 180Ω load is higher than 90Ω. With a 50Ω internal resistance, the maximum power transfer theorem states that the maximum output should be reached for a 50Ω load, but according to the tests, this maximum is obtained for a resistor of more than 90Ω. This is due to the fact that the theorem only takes into consideration ohmic losses, and not others, such as core and stray losses.

As seen in Figure 3.29, the force to move the rotor is fairly constant at about 28 lbf, during compression or rebound, indicating that it is predominantly governed by friction. The wavy behavior is the sign of the cogging forces, alternatively resisting and helping the motion. This accounts for variations of about 8 lbf, peak to peak. It can also be noted that the current drawn under the different loading conditions does not really impact the force required.

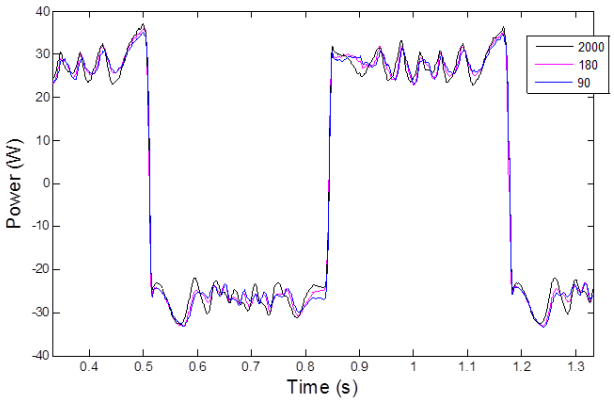

Figure 3.29: Forces under different loading conditions - ±0.75" at 1.5Hz

High friction prevents the system from being efficient and is one of the main reasons that led to only about 5 percent efficiency at most.

3.5.5 Correlation Simulation/Tests

With results from the tests on the actual prototype, it is then possible to determine how accurate the estimations of the output voltage were. There are several approximations in this process, including the quasi-static analysis of the magnetic field and the approximate of the derivative. The numerical model is slightly modified to take into account the differences that appeared during the machining process. The main difference is the inside diameter of the coil core, which is measured at 1.065 inch against 1.040 inch (initially planned). This effectively increases the air gap by over 50%, augmenting the reluctance of the magnetic path.

As mentioned before, the offset/starting point of the motion is particularly important in the shape of the output of the system. It is not possible to know exactly what that position was during the test, and it would be too time consuming and impractical to find it by trial and error, running a complete simulation for every attempt with the analysis in FEMM of all the positions, and then post-processing the data. So instead of assuming the exact motion before the magnetic analysis, the numerical model is tested along the entire range of possible displacements, yielding an exact knowledge of the normal magnetic flux inside the coil for positions at every 0.010in (Figure 3.30).

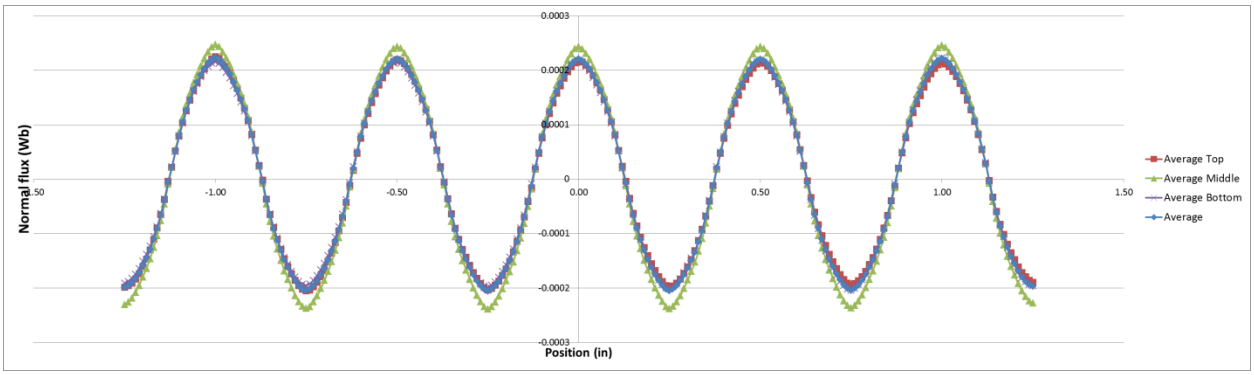

Figure 3.30: Normal flux inside the coil for every position

As can be expected, the normal flux presents a sinusoidal spatial distribution. It can be noted that the normal magnetic flux in the middle section of the coil is slightly higher than in the upper and lower sections. This can be attributed to leakage from the inner tube to the outer tube across the coil, bypassing the parts of the winding the furthest from the center.

Given an input of known amplitude and frequency, an offset, and a starting time, it is easy to determine the positions reached during the motion of the generator, and to which time and normal magnetic flux they are associated. With the magnetic flux linked to time, the next step is to compute the output voltage

by differentiating with respect to time, as stated in Faraday's law (see Section 2.5.3). The starting time is determined by the recorded data, and the offset is chosen to obtain the best fit possible.

The result (Figure 3.31) is a fairly good fit, but the estimated voltage from the simulation is quite wavy. Although the normal flux looks fairly close to an exact sine wave, some differences that can be imputed to the nature of the finite element analysis can be seen, and they lead to larger error after the differentiation with respect to time.

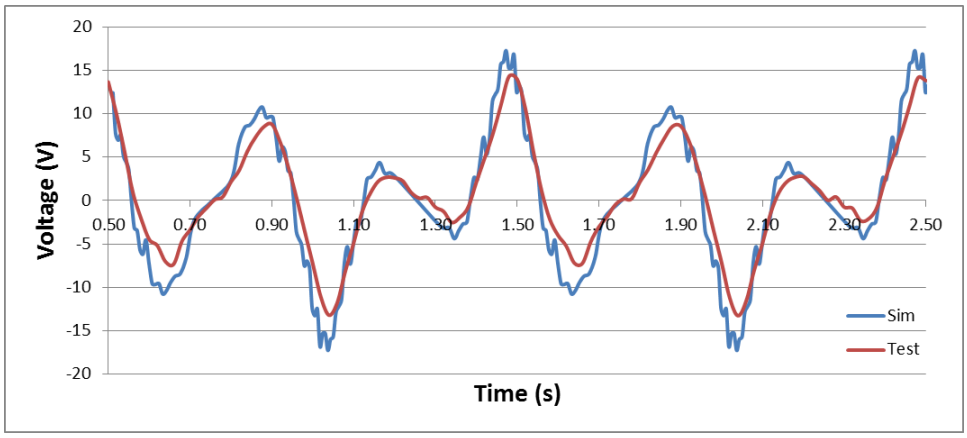

Figure 3.31: Voltage from the simulation and the actual test

One solution is to replace the normal flux from the simulations by an actual sine function. As seen in Figure 3.32, a sinusoidal model approximates quite well the actual output from the magnetic analysis.

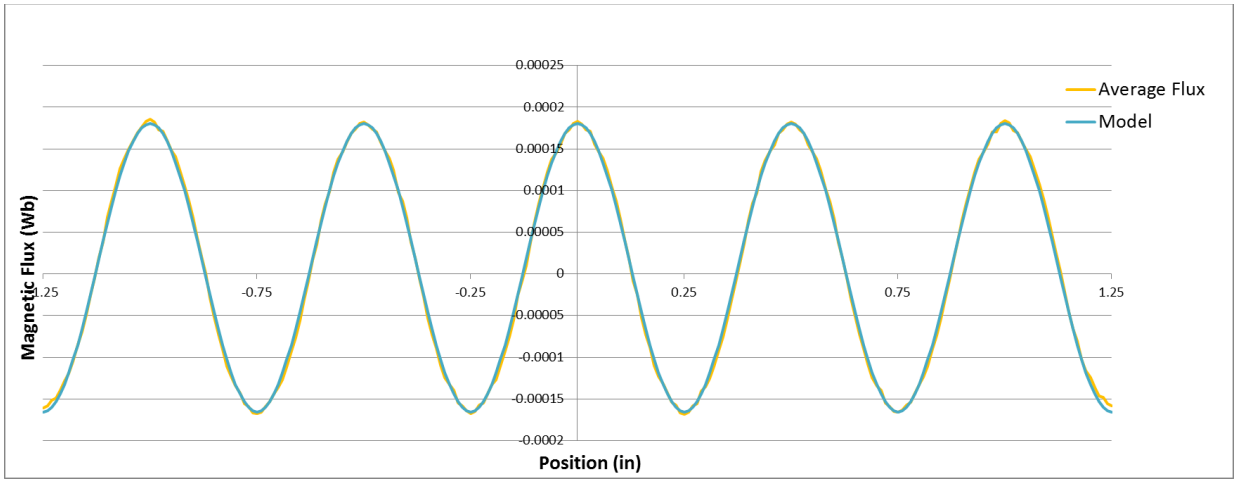

Figure 3.32: Normal flux and sinusoidal model

With the new model, the curve of the estimated voltage is a lot smoother, matching the result from the test of the actual prototype fairly well (Figure 3.33). Some small features are not represented exactly and the amplitude is slightly lower, but the overall shape is similar.

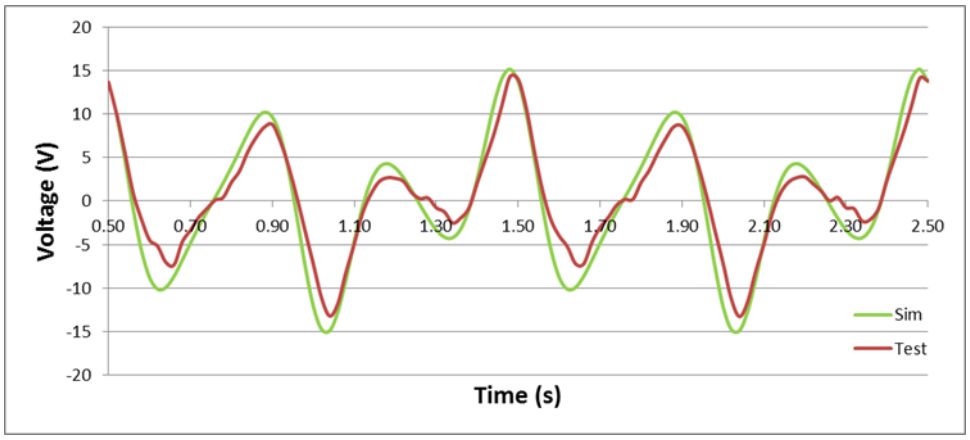

Figure 3.33: Voltage with the sinusoidal model of the magnetic flux and from the actual test

The voltages presented in Figure 3.31 and Figure 3.33 are from tests with an input of ±0.25 inch at 1Hz. The same process can be repeated to match any other test, for example ±0.75 inch at 1.5Hz, as on Figure 3.34.

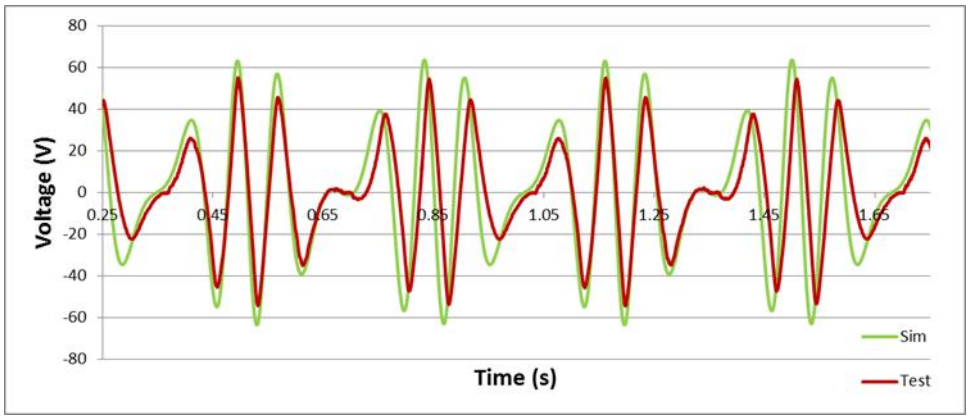

Figure 3.34: Voltage (simulated and from test) for a ±0.75in at 1Hz input

The general appearance of the simulated voltage is similar to that of the voltage from the test. The simulation slightly overestimates the output, and half of the cycle presents a timely fit, whereas the second half shows a shift. During the tests, some backlash was observed between the rod and the magnet arrangement. That play needs to be included in the model to get a better fit of the output voltage. So, at

the start of the rebound, the magnets don't actually move before the rod has moved a given amount. During the beginning of the compression, the gap is closed and then the magnet assembly starts to move. This model is illustrated in Figure 3.35 below, with a backlash of 0.050 inch.

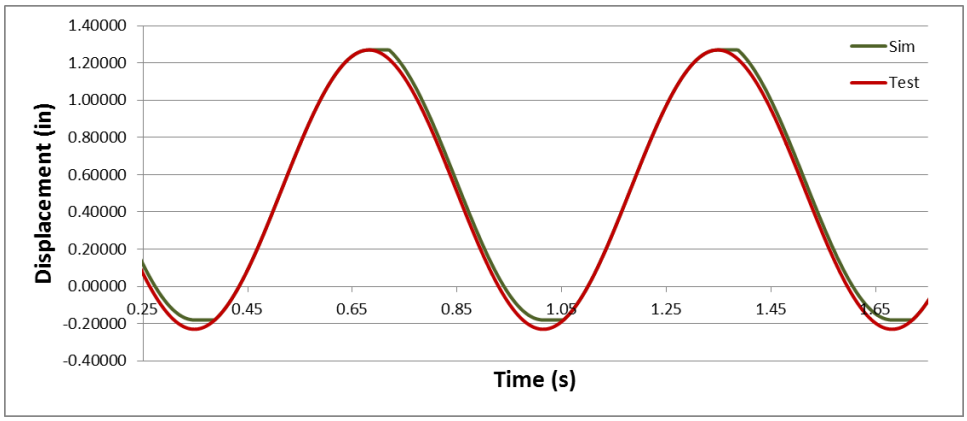

Figure 3.35: Displacement with backlash included

The result is a better match of the voltage recorded during the test (Figure 3.36). The curves are now in phase all along the cycles, while the voltages are still slightly overestimated.

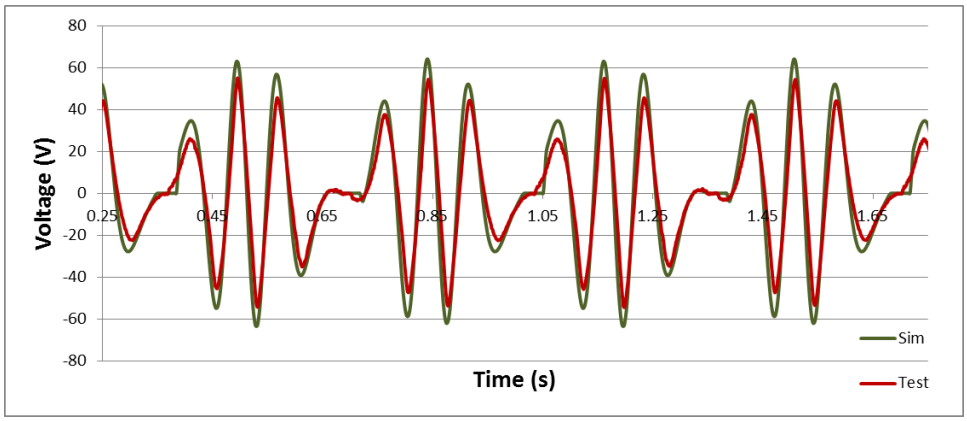

Figure 3.36: Voltage estimation including the rotor backlash

The main drawback of this backlash model is the flat portion visible between half cycles. In the actual prototype, some motion is still transmitted after the direction change, slipping toward the maximum difference. In order to replicate a behavior closer to that, the play is represented by a cosine function in addition to the totally free play (identical to what was used just before), now limited to 0.025 inch. The result is an improved match of the test output voltage, with a more limited flat line between changes of direction (Figure 3.37). The simulation is again slightly overestimating the output voltage (by about 12%). This can be attributed to different factors that are not accounted for. The magnetic path can be

more reluctant due to slightly different materials or due to small air gaps from machining imperfections. A simulation with gaps as small as 0.010 inch at the ends of the coil inner core will lead to an excellent match of the amplitude. The lower voltages could also be attributed to the presence of eddy currents. The magnetic path is not laminated and can be prone to significant eddy currents. The voltage from the test is also not a perfect sine wave, with some dissymmetry. This can be caused by some non-linear characteristics of the material, or some remanence that cannot be accounted for in the quasi-static model used here.

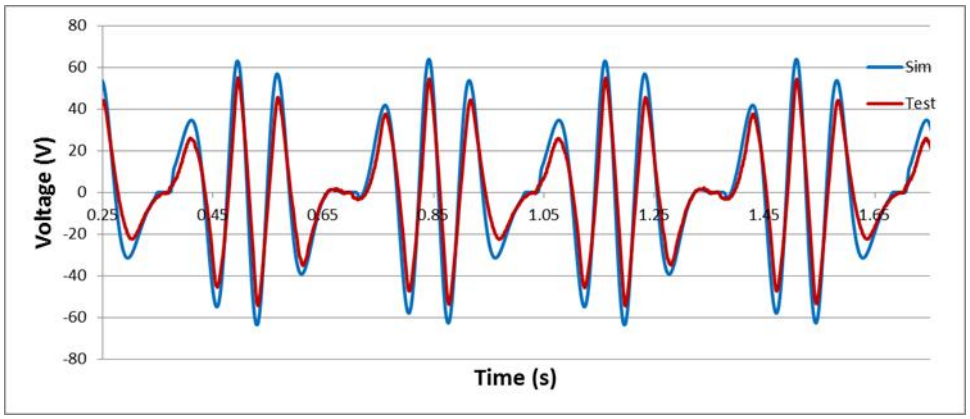

Figure 3.37: Voltage estimation with "improved" backlash model

Although the correlation between the voltages from tests and simulation for open circuits is good, when a load is connected to the circuit, the error becomes rather large. Since current is now flowing through the circuit, the internal resistance needs to be taken into consideration. The voltage across the load is then:

$$V_{load} = \frac{R_{load}}{R_{load} + R_i} V_{emf} \tag{3.4}$$

where R_i is the internal resistance of the generator. Figure 3.38 shows that, with an 180Ω load in the circuit, the simulated voltage drops by about 22%. But this still represents a 60% error compared to the voltage recorded during the tests. Including the impedance of the coil does not help reduce that error. So, the ohmic losses are not sufficient to explain the decrease of the output voltage. Other types of losses such as core losses with eddy currents, stray load losses due to the mmf of the load currents, and hysteresis losses need to be included in the simulation to offer an accurate estimation. Unfortunately, this is beyond the capabilities of the software FEMM, and other applications, like Maxwell 2D, should be used to accurately model the behavior of the generator and take into account all the transient phenomena that induce losses.

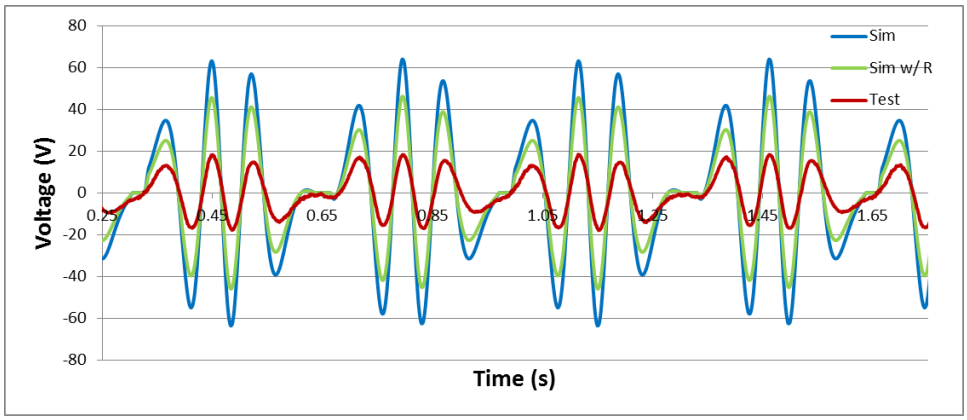

Figure 3.38: Voltages with an 180Ω resistor

The fact that an 180Ω load leads to more power generated than a 90Ω for a generator with a 50Ω internal resistance also highlights that there are more losses to consider than simply the ohmic losses. Otherwise, the maximum power would be reached for a 50Ω load according to the maximum power transfer theorem.

The software FEMM also has a function to estimate the forces due to magnetic attraction. However, since the model is perfect, with the rotor perfectly centered, there is no radial force and no friction force to estimate. Therefore, the only force that could be estimated is the cogging force.

3.5.6 Conclusion and Perspectives

The large coil design proved that it can generate enough voltage to potentially recharge a battery. Although the output voltages are good, the amount of power remains fairly low, far from the ultimate goal of 10-15 Watts. The losses (including ohmic) are too great for the prototype to be efficient. It would be difficult to eliminate the friction due to the attraction of the magnets to the inner tube. Inherently, a linear generator will always suffer from the need for a large number of magnets among which only a few are actually producing useful work; the others are only working against an efficient design.

The concept of a large coil could be improved, perhaps by splitting the large coil into two halves with or without a phase between them, relieving the areas with high magnetic field. It would also be possible to increase the diameter of the system to have larger magnetic path cross-sections.

3.6 Smaller Coils

In the design process, the ideas split into two paths: one leading to a very large coil with the idea that a large number of turns will lead to a useable voltage; and another remaining closer to multiple and more "conventional" smaller coils. The large coil proved that it is possible to generate a decent voltage even with small displacements, but it did not deliver in terms of power. The idea is then to test a system with a smaller coil in which the voltage generated will be lower but, hopefully, will allow for a lot more current to be drawn with lower losses, ultimately leading to high power being generated.

The chosen design (Figure 3.39) presents three coils which are one-half inch apart, so they don't use two adjacent magnets but skip one. This allows for larger coils (compared to a design with coils of the same width as a magnet) and it keeps them in phase.

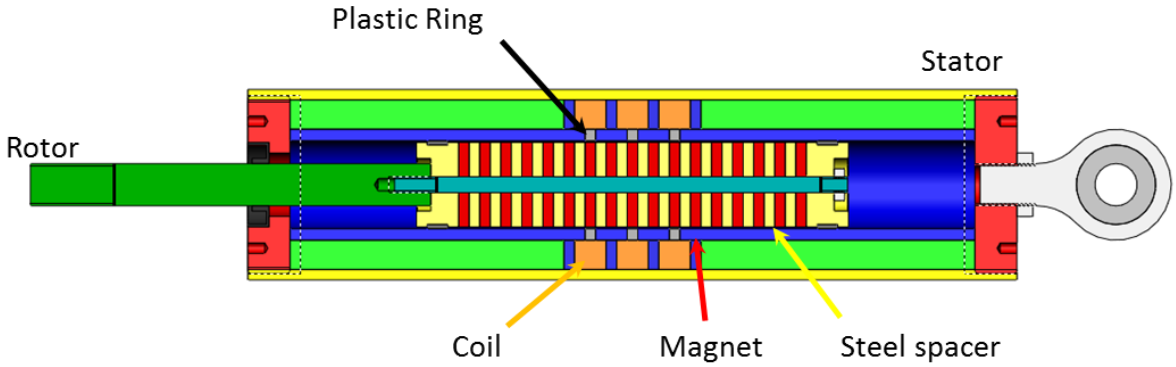

Figure 3.39: Small-coil design

The coil assembly was designed in such a way that all the other parts remain identical to those of the first prototype so they can be directly reused, avoiding the need for too many parts to be made for the new prototype. The new parts for the coil core were machined on our laboratory lathe. Once assembled, the lathe was used again to wind the coil, similarly to what has been done for the large coil of the previous prototype. Two coils were made identical with 225 turns of 24 AWG wire, and the third one was made with 350 turns of 26 AWG wire.

Once completed, the coils were kept in place with some tape (Figure 3.40) and the entire sub-assembly was slid into the outer tube. After the rotor was carefully put in place, the caps completed the new prototype.

Figure 3.40: New coil assembly

Unfortunately, one of the two identical coils became compromised during the assembly process and could not be used during the tests. However, this still left two good coils to conduct all the desired tests.

The Roehrig shock dynamometer was once again used to provide the input motion and the HP digital oscilloscope to record the output voltage across a resistive load. Various conditions (amplitudes, frequencies, resistors) were tested. The 350-turn coil alone only generates about one-half a volt for ±0.25"at 1Hz input (Table 7), far from the 6+ VRMS recorded with the previous prototype. The output powers are also lower but not with as much difference, with only a 50% decrease. It is important to note that the resistors used with this prototype are of significantly lower values: 50, 10, 1Ω. With the previous design, almost no power would be generated with such low resistance.

Table 7: Test results for the 350-turn coil

Input	Load	Coil - 355 turns	
		VRMS	Watts
±0.25"@ 1Hz	50Ω	0.615	0.008
	10Ω	0.455	0.021
	1Ω	0.244	0.060
±0.5"@ 1Hz	50Ω	1.177	0.028
	10Ω	0.841	0.071
	1Ω	0.361	0.130
±0.75"@ 1Hz	50Ω	0.485	0.005
	10Ω	1.227	0.151
	1Ω	1.735	3.010
±0.75"@ 1.5Hz	50Ω	2.578	0.133
	10Ω	1.790	0.320
	1Ω	0.593	0.352

With faster input velocities, the outputs evidently increase. With an amplitude of 0.75 inch and a frequency of 1.5Hz, the system can generate $2.6V_{RMS}$ with a 50Ω, resistor and over 0.3 Watts of average power.

The goal here is not to be able with only one coil to produce as much power as did the previous prototype. The advantage of small coils is that in the same space, several of them can be implemented. However, it has to be proven that the contributions of each coil can be added to reach larger outputs. The prototype is then tested with the two coils connected in series (Table 8).

Table 8: Results for each coil individually and connected in series

Input	Load	Coil 1 - 225 turns		Coil 2 - 350 turns		Both Coils	
		VRMS	Watts	VRMS	Watts	VRMS	Watts
±0.25"@ 1Hz	10Ω	0.391	0.015	0.367	0.013	0.766	0.059
±0.75"@ 1.5Hz	10Ω	1.549	0.240	1.725	0.298	2.396	0.574
		1.556	0.242	1.721	0.296	2.471	0.611
	20Ω					3.308	0.547
	50Ω	1.878	0.071	2.539	0.129	4.064	0.330
	100Ω					4.432	0.196
	10MΩ					4.841	0.000

We can note that the output voltage from the two coils in series is not equal to the sum of the voltages generated by each coil, for the same resistive load. For example, the coils in series produce "only" $4.064V_{RMS}$ when the sum of the individual contributions would be $4.42V_{RMS}$. That total is otherwise really close to the voltage generated with a resistor of twice the value, i.e. 100Ω. This can be explained by the internal resistance of the coils. Let's assume that each coil can be modeled as an ideal source of voltage in series with a resistor. The voltage generated by both coils in series across a resistor of value $2R$ is equal to twice the output from one coil across a resistor of value R. This can be easily clarified by the following equation:

$$V_{2R} = \frac{2R}{R_{\text{int},1} + R_{\text{int},2} + 2R}\left(V_{emf,1} + V_{emf,2}\right) = \frac{R}{R_{\text{int}} + R}\left(2 \cdot V_{emf}\right) = 2 \cdot \left(\frac{R}{R_{\text{int}} + R} V_{emf}\right) = 2 \cdot V_{R,coil,i} \quad (3.5)$$

assuming that $R_{\text{int},1} = R_{\text{int},2} = R_{\text{int}}$ and $V_{emf,1} = V_{emf,2} = V_{emf}$ and recognizing that the voltage from one coil across a resistor R is $V_{R,coil,i} = \left(\frac{R}{R_{\text{int},i} + R} V_{emf,i}\right)$.

The equation above assumes that both coils have the same internal resistance. However, the prototype has slightly different coils which leads to a slightly different expression:

$$\begin{aligned} V_{2R} &= \frac{2R}{R_{\text{int},1}+R_{\text{int},2}+2R}\left(V_{emf,1}+V_{emf,2}\right) \\ \Rightarrow V_{2R} &= \frac{R_{\text{int},1}+R}{\frac{R_{\text{int},1}+R_{\text{int},2}}{2}+R}\left(\frac{R}{R_{\text{int},1}+R}V_{emf,1}\right)+\frac{R_{\text{int},2}+R}{\frac{R_{\text{int},1}+R_{\text{int},2}}{2}+R}\left(\frac{R}{R_{\text{int},2}+R}V_{emf,2}\right) \end{aligned} \tag{3.6}$$

With the measured internal resistance, it yields:

$$V_{100\Omega} = 0.965\cdot V_{50,coil,1}+1.035\cdot V_{50,coil,2} \tag{3.7}$$

$$V_{20\Omega} = 0.870\cdot V_{10,coil,1}+1.130\cdot V_{10,coil,2} \tag{3.8}$$

The bigger the difference between the two coils and the closer those internal resistances are to the load resistance, the further away from 1 the two coefficients will be. This gives more dominance to the coil with the higher internal resistance.

The voltage from the coils in series across a resistor R will always be lower than the sum of the voltages from each coil, individually, across the same resistor R. This can be verified with the following derivation:

$$\begin{aligned} V_R &= \frac{R}{R_{\text{int},1}+R_{\text{int},2}+R}\left(V_{emf,1}+V_{emf,2}\right) \\ \Rightarrow V_R &= \frac{R_{\text{int},1}+R}{R_{\text{int},1}+R_{\text{int},2}+R}\left(\frac{R}{R_{\text{int},1}+R}V_{emf,1}\right)+\frac{R_{\text{int},2}+R}{R_{\text{int},1}+R_{\text{int},2}+R}\left(\frac{R}{R_{\text{int},2}+R}V_{emf,2}\right) \\ \Rightarrow V_R &= \frac{R_{\text{int},1}+R}{R_{\text{int},1}+R_{\text{int},2}+R}V_{R,coil,1}+\frac{R_{\text{int},2}+R}{R_{\text{int},1}+R_{\text{int},2}+R}V_{R,coil,2} \end{aligned} \tag{3.9}$$

Using the measured internal resistance, it leads to:

$$V_{50\Omega} = 0.890\cdot V_{50,coil,1}+0.954\cdot V_{50,coil,2} \tag{3.10}$$

$$V_{10\Omega} = 0.661\cdot V_{10,coil,1}+0.859\cdot V_{10,coil,2} \tag{3.11}$$

Using the results from the tests, it can be verified that, given the voltages generated by each coil across a 50Ω resistor, these equations offer a perfect fit to the output voltages of the coils in series, both across a 100Ω and a 50Ω resistive load (Figure 3.41).

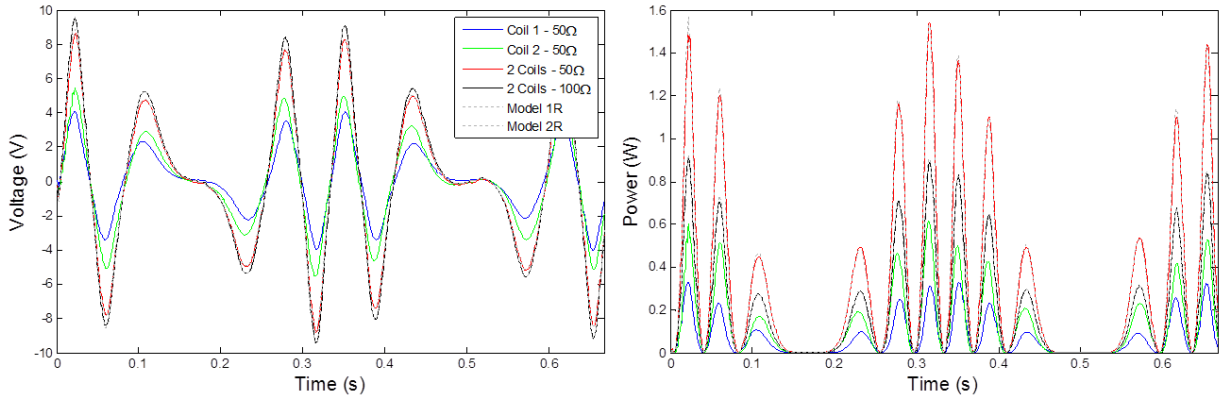

Figure 3.41: Small coil prototype - Outputs with 50Ω and 100Ω loads

The equations still give a good approximation for 10Ω and 20Ω loads (Figure 3.42). The discrepancies between the model and the actual test measurements can be attributed to the fact that the assumption that the coil can be modeled as an ideal source of voltage and a resistor in series is no longer totally valid.

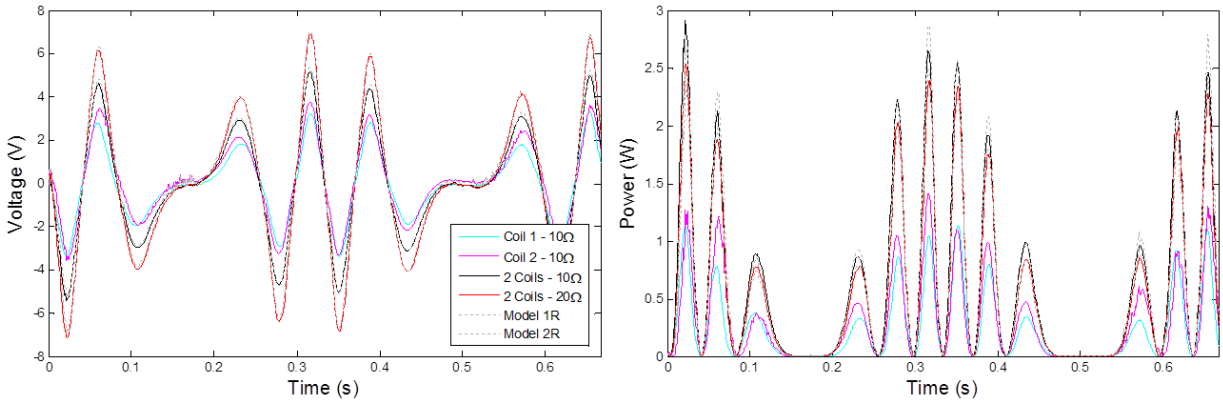

Figure 3.42: Small coil prototype - Outputs with 10Ω and 20Ω loads

If this simple voltage source/resistor model holds, the output voltage across a resistor R is defined by:

$$V_R = \frac{R}{R_{\text{int}} + R} \cdot V_{open-circuit} \tag{3.12}$$

This is the equation of the blue curve in Figure 3.43. It can be seen that the test results match that curve quite well above 40Ω. Below 40Ω, the voltages are lower than expected, and taking into account only ohmic losses is no longer sufficient. This explains why the equations give very good results for the 50Ω tests, and yields to slight differences with the 10Ω load. It is worth noting that the error is limited because the voltage from the two coils in series is estimated using the voltages of each coil across the same resistor (10Ω). Using the open-circuit voltage would have led to a significantly larger error. The fit of the curve for 50Ω and 100Ω, however, should still be extremely good, as the measured voltage is really close to the expected value (Figure 3.43).

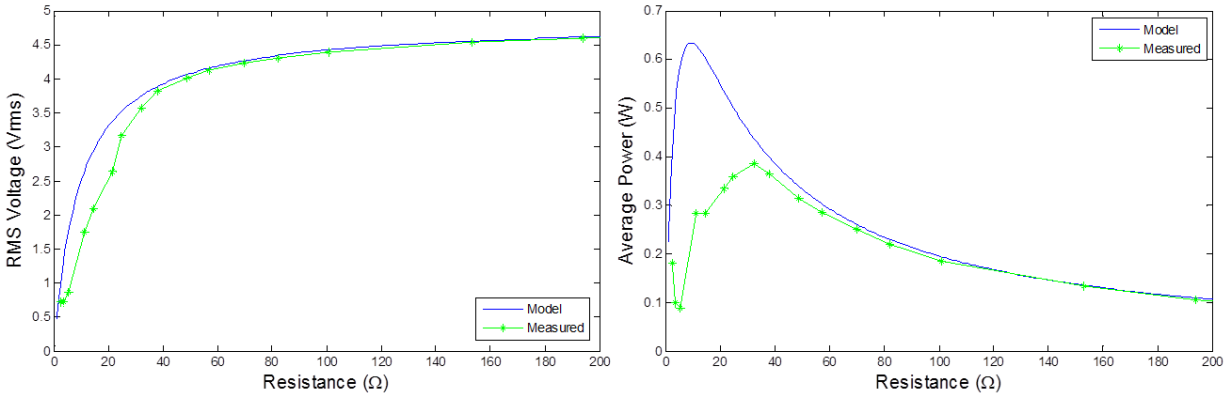

Figure 3.43: Voltage and power of the two coils in series - ±0.75in at 1.5Hz

Friction is still an important part of the force required to move the system (Figure 3.44) and it remains in the same range of values: around 35lbs. But, unlike in the previous prototype, there is no plastic ring to counterbalance the gap between the stator poles. So, the cogging forces are significantly larger than before, generating large variations of 50 lbs, peak to peak. The efficiency is comparable to that of the previous design, limited to only a few percent.

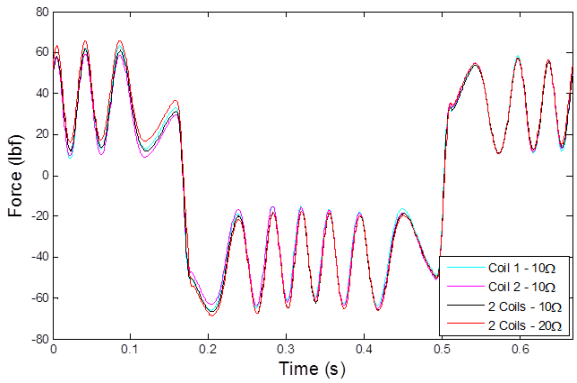

Figure 3.44: Forces under different conditions - ±0.75" at 1.5Hz

3.6.1 Concluding Remarks

The prototype using small coils, as expected, provides lower voltages, but the lower internal resistance and the lower overall losses allow higher currents, which in turn, translate to power. With only two coils, it is possible to obtain slightly more power. Furthermore, in the same given space, it is possible to add more coils and improve the performance and the power generated.

3.7 Conclusion

The two linear generator prototypes have proven that it is possible to generate, even from a small input (±0.25in at 1Hz), a voltage high enough to be useful and to potentially charge a battery. In terms of power, the prototypes were able to provide up to an average of 0.6 Watts or 3 Watts at peak. This may be too limited for some applications, but there is the possibility of increasing those values by increasing the number of coils. An important drawback of the system is that the magnets generate significant friction between the rotor and the stator, leading to a low overall efficiency (below 5%).

Chapter 4

Rotating Generator Energy Harvesting Systems

4.1 Introduction

The idea behind the design of a linear generator was to keep the system very simple and avoid the losses that are inherently present in mechanical systems, such as in a translation-to-rotation transformation. However, this comes with limitations, especially with the small amplitudes (less than 1 inch) and low frequencies (between 1 and 2Hz) that are considered here. Furthermore, as explained by Faraday's law, variations in the magnetic flux are needed to produce an electromotive force. In a linear generator such as that presented in the previous chapter, there is no solution for a given input to increase the speed or frequency of these variations. But with a rotating generator, it is easy to place a gearbox to increase the motion. This implies mechanical losses in the process, but it allows the generator to then be in a range of speed where the electromagnetic induction is significantly more efficient. The overall process would then be far more efficient. This might imply higher forces, but this is not a limitation (up to a certain extent) for a rail car suspension.

This chapter presents the work that has been done to develop an energy harvesting system using a rotating generator, after the transformation of the linear motion of the suspension into a rotation.

4.2 Preliminary Prototype

4.2.1 DC Motor Testing

For the preliminary series of tests, small DC motors are used. They are very easy to find and quite inexpensive. The exact description of the chosen motor is a 9 to 18VDC Hobby Motor from Radioshack (Model: 273-256). It can work with voltages up to 18V and current up to 2.4 Amps. The maximum speed is 24000rpm without load and 18000rpm with load. They measure about 1 inch in diameter and 1.5 inch in length (not including the shaft). Therefore, they are quite compact and can take over 40 watts.

Prior to building the system, the motor was first tested alone. The setup is two motors with their shafts linked together (Figure 4.1): one is used as a motor and the other is used as a generator, from which the voltage is measured. This allows us to determine the voltage, the power they can produce, and at what kind of speed.

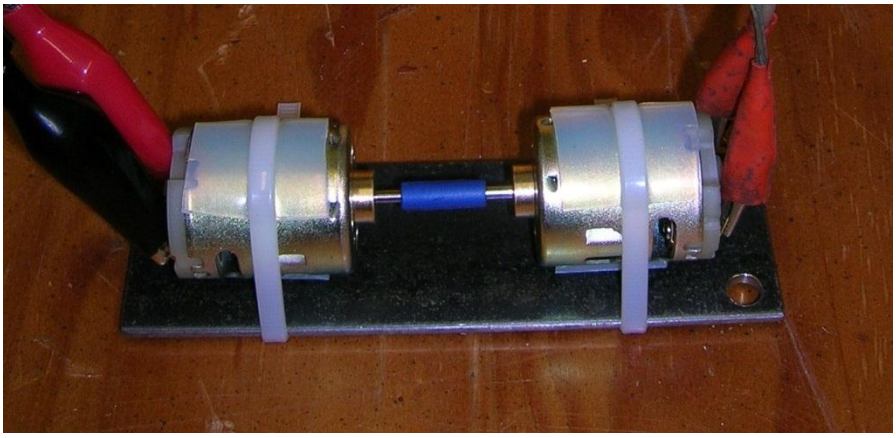

Figure 4.1: Motor test setup

Not surprisingly, the higher the speed, the higher the voltage, which can easily reach over 12 Volts. The motors were not pushed to their limits or tested in hard conditions. The tests are really to get an idea of the range of possibilities, keeping in mind that the speed will have to be obtained by a rack and pinion/gearing system from a relatively small initial displacement. To estimate the power the voltage across a 10Ω resistor was measured with a voltmeter. Over 5 Watts were easily obtained.

It is also worth comparing the energy going in and that was recovered. From the measurements, a rough estimation of the efficiency is around 40 to 50%. So, less than half of the supplied energy is recovered at the output.

The conclusion of this small preliminary test is that with moderate speed, it is possible to obtain reasonable output power. Now we will attempt to determine what can be achieved with the rack and pinions and the gearing system.

4.2.2 First Quick Prototype

The main idea in this prototype (Figure 4.2) was to quickly and inexpensively estimate the potential of a system using a rotating generator. The generator was the DC motor presented previously. All the other mechanical parts were repurposed elements from toys: the rack and pinion was from a fly wheel toy, with

a rack that is about 10 inches long with a handle at one end. Its pitch is about 8. The matching pinion has 10 teeth.

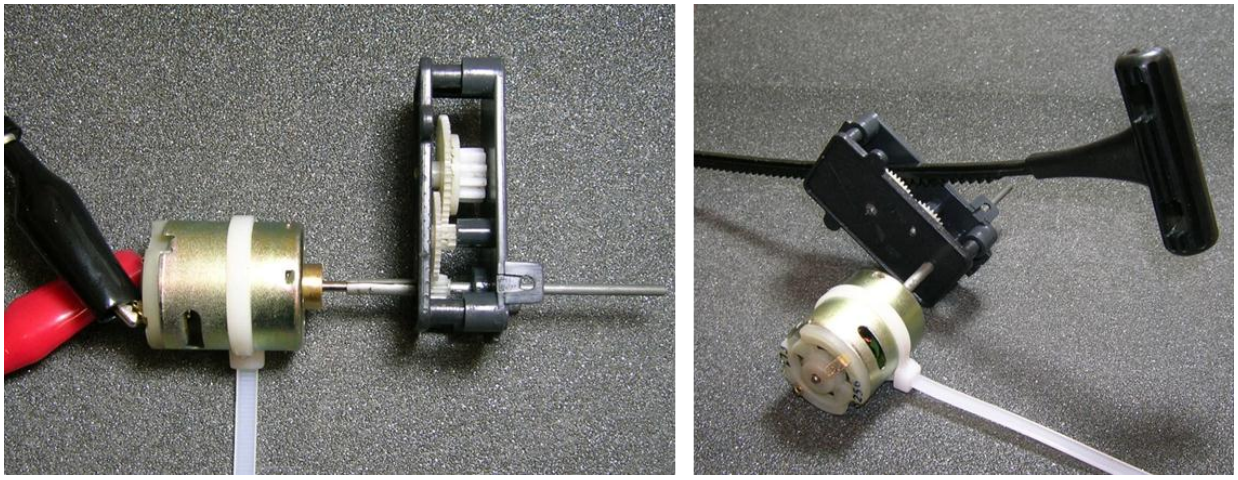

Figure 4.2: First quick prototype

The gearing system comes from a spring-loaded car, with gears rearranged to be more suitable for the purpose of the project. The spring was, of course, removed. The first gear has 50 teeth and drives a 9-tooth pinion which is on the same axis as a 36-tooth pinion. The final gear has 10 teeth. This leads to a 20:1 ratio gearbox. The gearbox comes with a sort of freewheel system. The shaft of the intermediate gears is in an oblong slot, allowing the loss of contact between gears in one direction. Originally, this allowed the spring to drive the wheels of the car; when the spring is fully unwound, the car can keep going without fighting against the spring. This feature allows us to reset the rack without making the motor turn in the opposite direction. It is not efficient, with a lot of backlash, but it turns out to be a very helpful feature.

The prototype is certainly not perfect and it is probably not very efficient. However, that is not what is discussed here; it will be a concern in a more well-thought design.

4.2.3 Testing

The electric motor (used as a generator) was connected to a resistor. The voltage across that resistor was then recorded using a digital oscilloscope. The output power was then derived using Ohm's law.

Two different resistors were tested: 10Ω and 1Ω (both rated for 10W). About 7 inches of the rack were used during the test with the 10Ω resistor. Since the rack is pulled by hand, the input is not known

exactly. The velocity could be estimated with the recording of the voltage where it is easy to see the time of traction and reset. It is not really important, though, to determine the actual speed of the rack at that point. The rack obviously moved faster than for a realistic simulation: a fraction of an inch of amplitude with frequencies between 1 and 2Hz. But that means that in the actual prototype, the gear ratio or the generator will be chosen to compensate.

The tests (Figure 4.3) show that up to 6 Volts at peak can be produced, yielding 3.5 Watts of power. On average, the output voltage is around $3V_{RMS}$ and the average power is 0.9 Watts.

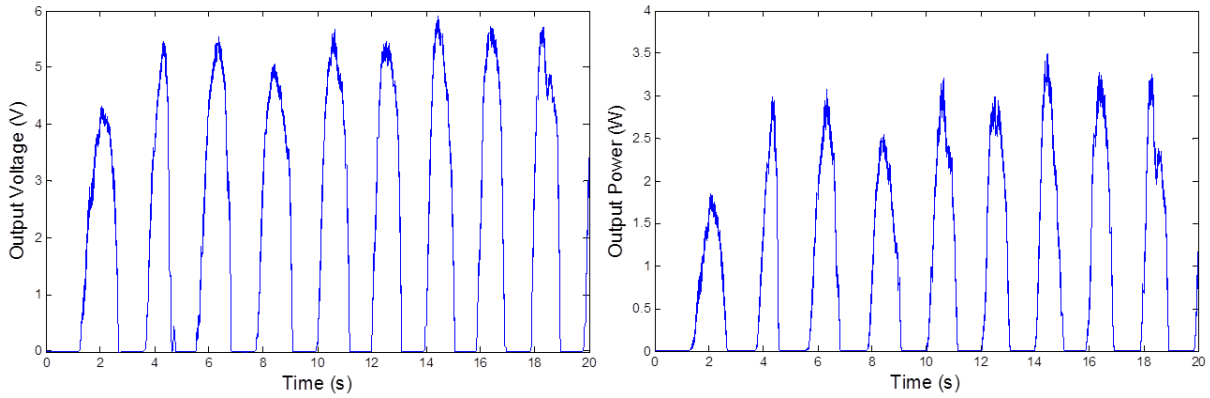

Figure 4.3: Quick prototype - Test results with a 10Ω resistor

For the test with the 1Ω resistor, more force is required to pull the rack. This leads to a shorter portion of the stroke being used, thereby limiting the speed of the rack. This partially explains the lower output voltage and power. The voltage peaks at 1.5 Volts and the power at 2.3 Watts. The average values are 0.8 V_{RMS} for the voltage, and 0.65 Watts for the output power.

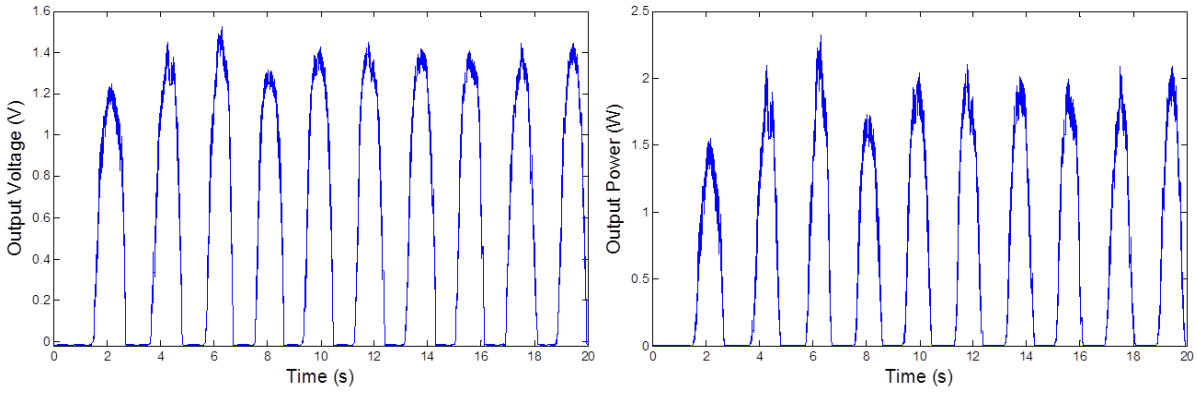

Figure 4.4: Quick prototype - Test results with a 1Ω resistor

4.2.4 Conclusion

The test results of that quick prototype are good and quite encouraging. It has to be reiterated that only one direction is used here. This means that only half of the displacement is used. It would be not accurate to just multiply the average values by 2, since the reset time is a little shorter than the work time, but it would give an idea of the maximum values to expect. It also has to be considered that by using both direction motions, the velocity of the motor would not go to zero after each stroke, and would require less effort to make the generator turn, thereby likely making the output power more stable.

As the prototype is currently designed, the test shows the possibility of generating around 3 V_{RMS} and an average power at almost one Watt. By comparison with the values obtained with the linear generator, the power generated with the rotary generator is much better: the output power is 50% higher with the one-way motion, and it could be more than double using up and down displacements. One aspect to highlight is that for all the tests, the rack was pulled by hand. This means that the force required here is not really high, and certainly lower than that for the linear generator.

The design of the linear generator still has room for improvement, but the latest tests tend to prove that there is a lot more potential in using the rotating generator, coupled with a gearbox used to increase the speed.

The next step is to bring the concept of using a rotary generator with mechanical transformations to a damper-like prototype that could be tested in a repeatable and controlled environment, like on the Roehrig shock dynamometer. This would allow a fair comparison between the rotating generator and the linear generator.

4.3 First Generation Prototype

4.3.1 Initial Design

4.3.1.1 *Concept*

The promising results from the simple pulled-by-hand rack-and-pinion prototype led to the development of a practical damper-like energy harvester. One of the first aspects of the design was to decide how to transform the linear motion of the suspension into the needed rotation. The choice was between two mechanical devices: a rack and pinion mechanism, similar to what was used previously, or a ball screw.

The ball screw is a mechanism primarily used to transform rotation into linear motion, but due to very low friction, it is reversible and it can very well be back-driven. The rack and pinion presents a challenge due to the axis of rotation of the pinion being perpendicular to the general axis of the system. On the other hand, the ball screw generates a rotation that can more easily be exploited. In this first design using a rotational generator (Figure 4.5), the rack and pinion is then replaced by a ball screw to transform the translation of the suspension into a rotational motion. A shaft then connects the ball nut, on the very left side, to a planetary gearbox, on the other side of the harvester. The use of a harmonic drive was first investigated. Compactness, light weight, the possibility of high gear ratio, and no backlash are a few aspects of the mechanism that make it an attractive option. However, this type of drive is primarily intended to reduce speed and not increase it. The starting torque to back-drive the system, almost by itself, prevents its implementation. A planetary gearbox is ideal for the present application. It is rather compact, allows large gear ratio, can be driven to increase speed, has coincident input and output axis of rotation, and its shape fits very well in a damper-like package. The output of the gearbox is then connected to the generator, at the far right side of the prototype.

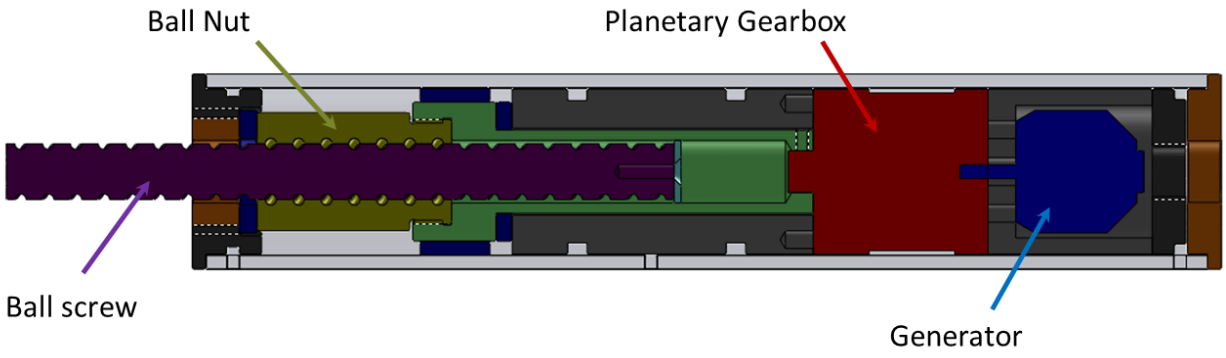

Figure 4.5: Design of the first prototype using rotary generator

4.3.1.2 *Prototype*

With the design defined, the prototype can then be built. Unlike the linear generator where every part has to be custom made, the new prototype can extensively rely on off-the-shelf parts. The ball screw was chosen among the standard dimensions readily available and was ordered from one of many suppliers. After some research, a planetary gearbox that had the right dimensions to fit inside the system was found at a robotic part supplier, Banebots. The standard ratio for the gearbox was chosen to be 20:1, in relation to the characteristic of the generator. However, parts can be interchanged and the gear ratio adjusted if necessary.

A permanent magnet DC brushless motor was used as a generator: an Alpha 400 from Exceed RC (Figure 4.6). It is a motor originally designed for RC airplane applications. The two main criteria were its size, 1.1in diameter, and its Kv value. The Kv is the ratio of the speed of rotation over the voltage, always expressed in rpm per volt. A low Kv is desirable so that higher voltages can be reached with lower speeds. It is inversely proportional to the back-emf constant K_e, usually expressed in Volts per rad/s.

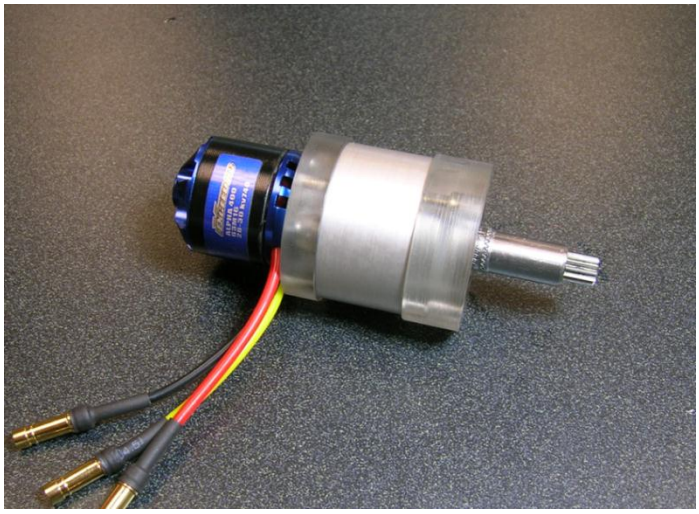

Figure 4.6: Exceed RC Alpha 400 motor mounted on the planetary gearbox

The motor chosen here has a Kv of 740rpm/V, meaning that every 740 rpm yields 1 volt and so, to reach 12 volts, a speed of 8880 rpm is required. This really highlights how important the speed is and the benefit of using a gearbox, one of the main advantages of this design over a linear generator where it would be impossible.

It can be noted that the motor is an outrunner. This means that the coils, on the stator, are inside, and the magnets on the outside magnet ring are rotating around it. By contrast, an inrunner has the rotor and its magnets inside, with the coils around it. Their Kv is usually significantly higher than outrunners, but they have lower inertias. They are mainly used for RC cars, whereas outrunners are primarily found in airplane and helicopter models.

Although many parts are standard, few parts had to be specifically made. The tube, spacers, and the end caps were machined in the laboratory. The choice for material went to polycarbonate, which is strong enough to withstand the forces experienced by the system, while offering a clear view to what is happening inside. The shaft, more complex with atypical inside threads, had to be outsourced.

With all the parts on hand, the prototype was fully assembled (Figure 4.7) and ready for testing.

Figure 4.7: First prototype mounted in the Roehrig shock dynamometer

4.3.1.3 *Testing*

The tests were once again conducted on the laboratory's electromagnetic actuated (EMA) shock dynamometer from Roehrig Engineering. For all the tests, the input was a pure sine wave, with a 1 Hz frequency. Three different amplitudes were tested: 0.25in, 0.5in and 0.75in. For each input, four different loads were tested: open circuit, i.e. no load beside the 10MΩ of the oscilloscope, 10Ω, 2.5Ω, and 1Ω. The output voltage is rectified through a 6-diode bridge before being fed to the resistive load. The voltage across that resistor is measured and recorded on a digital oscilloscope. The output power is then computed, during post-processing in Matlab, using Ohm's law, and following the formula:

$$P_{out} = \frac{V_{out}^{\ 2}}{R_{load}}$$

Although a little bit on the low side, the voltages show good improvement over the output of the linear generator, especially considering the voltage drop due to the diodes (about 1.4 Volts). In open circuit, the

voltages range from 7V to 12V, and for the lower resistance, 1Ω (Figure 4.8), they are down to as low as 1.1V or 5.5V depending on the input. Although the improvements in terms of voltages are not extremely impressive, the output powers show an extremely large progress compared to what is produced by the linear generators. It is now possible to produce an average power of more than 1.3 Watts with only 0.25in at 1Hz of input motion, and it can reach as high as 30 Watts with 0.75in at 1Hz. At peak, the power reaches close to 5 Watts, with the smallest input to about 68 Watts with larger displacements.

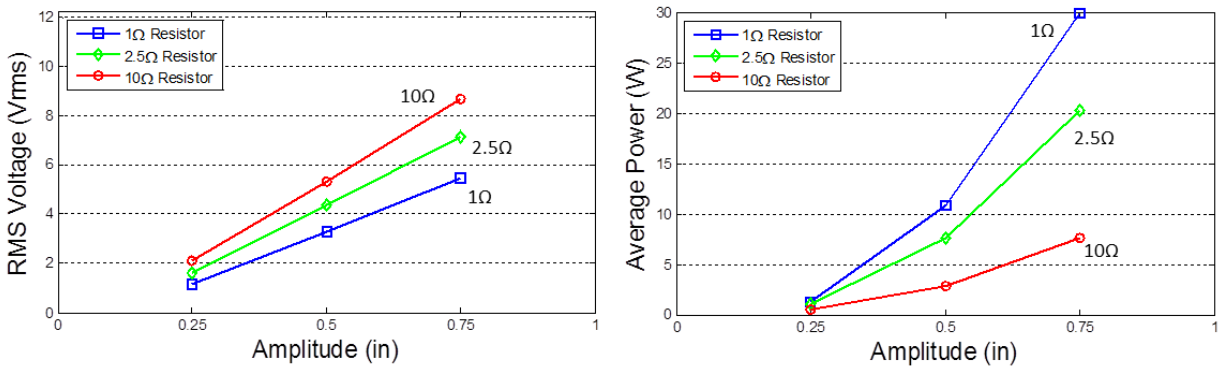

Figure 4.8: First prototype - Tests at 1Hz

Considering the results obtained with the 10Ω resistor, the power from this new prototype is about 10-fold higher than that from the linear design. Whereas the 1Ω resistor does not help to increase the output power with the linear generator, it makes the output power quadruple (from 10Ω) with the rotational prototype.

With a rotational generator, all the magnets contribute to create the output power, unlike in the linear design, where only two magnets at a time work. This makes a significant difference in the performance. Again, the speed of variation of the magnetic field is important for the output voltage. So, with the gearbox, the second design has a great advantage.

In terms of force, the system can now require up to a few hundred pounds for the fastest tests. For train suspension applications, this is neither an issue nor a limitation. A lot more force is required but with the very large increase in output power, the overall efficiency is rather good, being reportedly between 15% and 25% for the slowest input, and between 34% and 45% for the fastest.

One of the main drawbacks of this design is that with the ball nut rigidly connected to the gearbox input shaft, the generator has to come to a complete stop, between each oscillation, before starting to rotate in the opposite direction. As the generator stops, the voltage drops to zero (Figure 4.9).

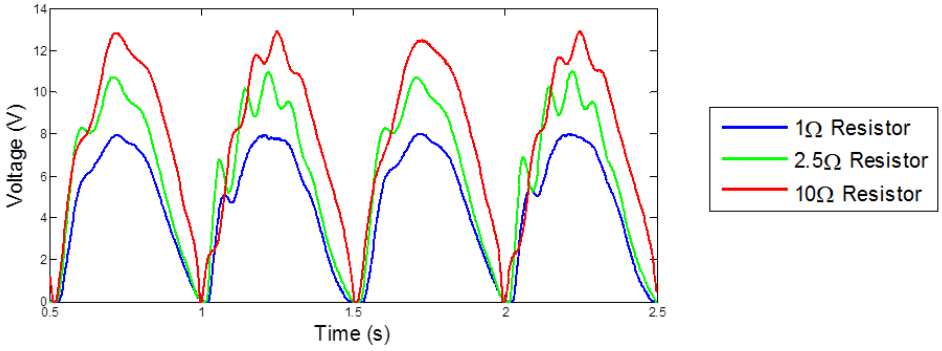

Figure 4.9: Output voltage – first prototype – ±0.75in at 1Hz

This prevents the voltage from being smooth. On the mechanical side, the inertia of the rotating parts cannot be kept and, worse, it must be fought during every cycle. This means that some unnecessary forces are required to first slow down the parts in rotation, and then more forces are needed to reaccelerate them. A solution is to implement a system that leads to the generator always rotating in the same direction, and that allows it to spin freely and to keep turning on its inertia when the input is slowing down. A solution to achieve all that is the upgrade presented in the next section.

4.3.1.4 *Conclusion*

The first damper-like prototype using a rotational generator does not disappoint on the promises of the preliminary tests: it allows fair voltages to be generated and significant amounts of power. It also performs quite efficiently, with up to 44% of the input energy dissipated in the test resistor. The rigid connection between the ball nut and the gearbox constraining the generator to move back and forth is an aspect of the prototype to be improved upon in order to further enhance already good performance.

4.3.2 First Upgrade: Keeping the Momentum

4.3.2.1 *Concept*

To improve the performance of the energy harvester, clutch bearings are incorporated in the design in such a way that the generator can keep spinning freely with its inertia if the input is slower, and such that it is always rotating in the same direction. Clutch bearings, also called one-way bearings, are a type of freewheel mechanism. They disengage the driven shaft from the driving shaft when the latter rotates slower. They also allow torque to be transmitted only in one direction. This is the same principle as what

is found in the freewheel of a bicycle where the pedals can drive the wheel but cannot be driven by it. The clutch bearing used for the prototype looks like regular needle bearings. However, the needles can slide in small slots: in one direction they slip, while in the other, they lock between the shaft and the bearing housing.

It is easy to understand how clutch bearing would allow the generator to keep rotating on its own inertia without being slowed down by the decelerating ball nut. The real challenge is to find a way to make both up and down motions create a rotation in the same direction. Systems using bevel gears have been investigated, but the available space is too limited to implement such a mechanism. The solution is to ingeniously use the first stage of the planetary gear. The elegant yet simple solution comes from the realization that whether the ring gear rotates clockwise or the planet carrier is turning counter-clockwise, the sun gear will rotate in the same direction: counter-clockwise (Figure 4.10).

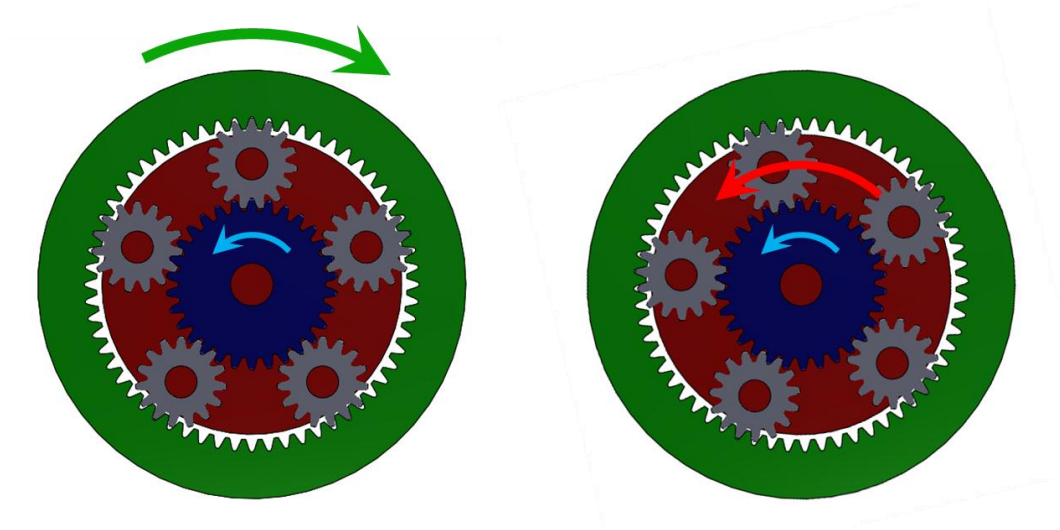

Figure 4.10: Planetary gear motions

The ball nut is then connected, through clutch bearings, to the ring gear and the planet carrier, transmitting power to one or the other depending on the direction of the rotation.

It is also important to prevent the ring gear and planet carrier from rotating backward. Otherwise, when the ring gear rotates clockwise, since there is a resistive torque due to the generator on the sun gear, the planet carrier would be forced to rotate clockwise with the ring gear. Then no power is transmitted to the generator. Also, when the planet gear rotates counter-clockwise, the ring gear would be forced to turn in the same direction, locking itself against the ball nut and leading to a 1:1 overall gear ratio. So, two

additional clutch bearings are required to prevent the ring gear and the planet carrier from rotating in the wrong direction and to obtain the desired gear ratio. This brings to four the number of one-way bearings needed.

The two inputs and the output have the same axis of rotation, making that solution well-suited to fit inside a damper-like device. The main drawback is that the gear ratio is different on the two power paths. When the planet carrier is the input and the ring gear is fixed, the gear ratio is:

$$\frac{\omega_{SG}}{\omega_{PH}} = \frac{2\left(N_{SG} + N_{PG}\right)}{N_{SG}} = 1 + \frac{N_{RG}}{N_{SG}} \tag{4.1}$$

where ω is the speed of rotation, N is the number of teeth, and the subscripts SG, PG, and RG refer to the sun gear, the planet gear, and the ring gear, respectively.

When the ring gear becomes the input, the gear ratio transforms to:

$$\frac{\omega_{SG}}{\omega_{RG}} = -\frac{N_{RG}}{N_{SG}} \tag{4.2}$$

So, the gear ratio will inevitably have a difference of 1. For example, with the gear used in the prototype, the gear ratio will be 5:1 when the planet carrier is the input, whereas it is only 4:1 when the ring gear is driving the gearbox. This is not ideal, but it is still a good trade-off to have the generator always rotating in the same direction and no longer being brought to a complete stop between oscillations. It is worth mentioning that shock absorbers usually have different settings in rebound and compression. So, it might be desirable to have two different levels of resistive force depending on the direction of travel of the suspension.

The upgraded prototype (Figure 4.11) was designed to reuse as many parts as possible to limit to a minimum the parts that need to be machined. The ball screw/nut, the shaft (that could not be machined on the laboratory's lathe), the containing tube and its caps, the gearbox adapter, and the generator were those used previously. The new parts are located in the mid-section of the system. The ring gear is no longer fixed to the outer tube, but is now attached to a hollow shaft that can be driven by the main shaft rigidly connected to the ball nut. The planet carrier also adopts a different shape to be able to receive a clutch bearing. It is also elongated across the planet gears to offer a possible bearing surface for a clutch bearing attached to the tube. This is the smart trick to allow contact between the fixed tube and the planet carrier, otherwise fully contained inside the ring without any access to non-rotating parts.

The new parts need a little more space lengthwise, but fit nicely in the tube inner diameter.

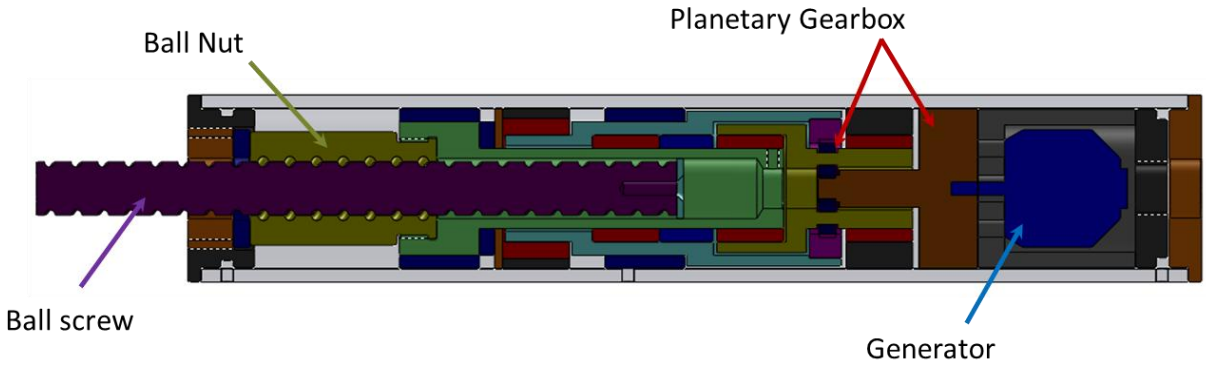

Figure 4.11: First prototype upgraded with clutch bearings

4.3.2.2 *Prototype*

The new parts, primarily consisting of the ring gear shaft and the modified planet holder, are machined in the laboratory, again using polycarbonate. This choice of material is useful to have a clear view of the internal moving parts of the system, but is not the best in terms of power transmission. Polycarbonate is too soft to lock the needles of the clutch bearings. The bearing surfaces were then replaced by aluminum parts, solving the slipping issue.

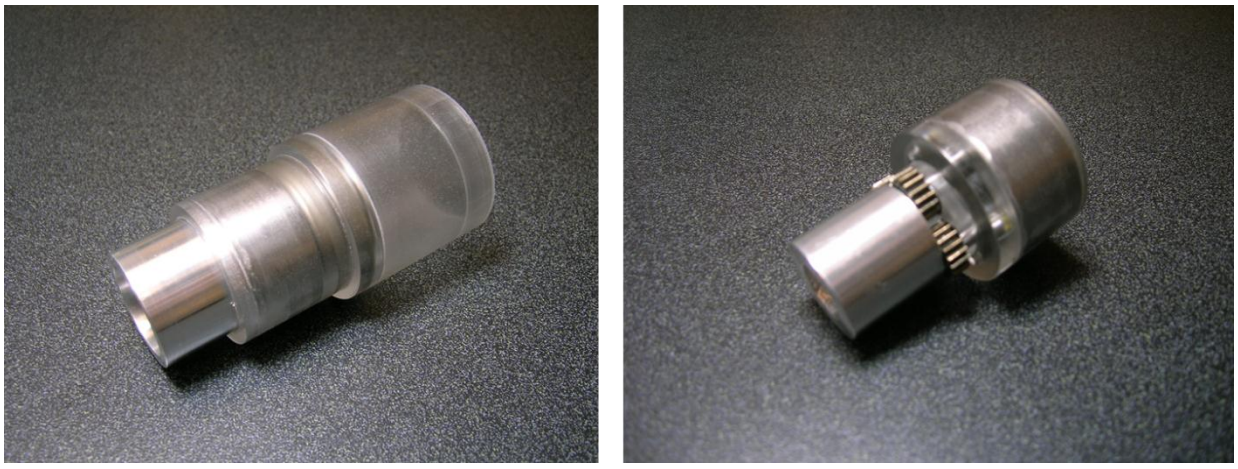

Figure 4.12: Ring gear shaft (left) and modified planet holder (right)

The new parts were paired with one-way and regular bearings, and the entire system was put back together and ready for testing.

4.3.2.3 *Testing*

In Figure 4.13, it is easy to see that all the voltages are not forced to drop to zero anymore. For instance, the open-circuit voltage now varies between 8V and 25V. It can be noted that, with low resistance, the voltage can effectively drop to zero. This is due to the resistive torque of the generator, which is strong enough to decelerate the rotating parts down to a complete stop before the ball nut has time to get back to speed on the next oscillation.

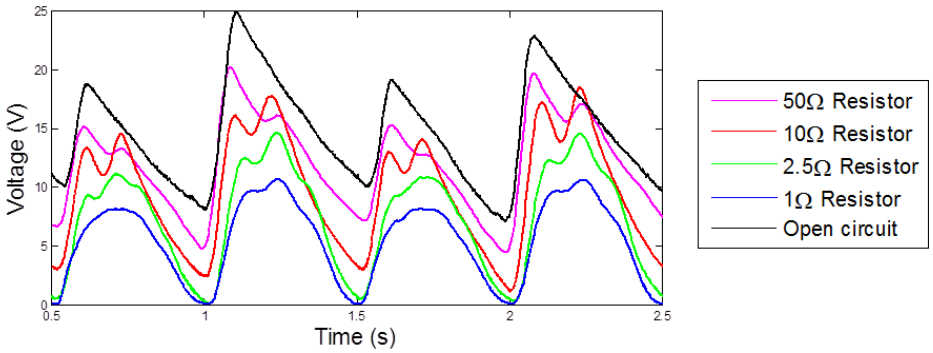

Figure 4.13: Output voltage – prototype with clutch bearings – ±0.75in at 1Hz

With the modifications of the clutch bearings, the output open-circuit voltage now reaches up to 15.5V_{RMS} and the average output power is recorded over 40 Watts for one test. This shows an average increase of about 25% in voltage, and about 50% in power. The increases are larger for larger resistance, where the inertia can provide a significant difference in the velocity of the generator. For very small amplitudes, the output remains fairly low, with voltages of only a few volts and only a few Watts of power produced.

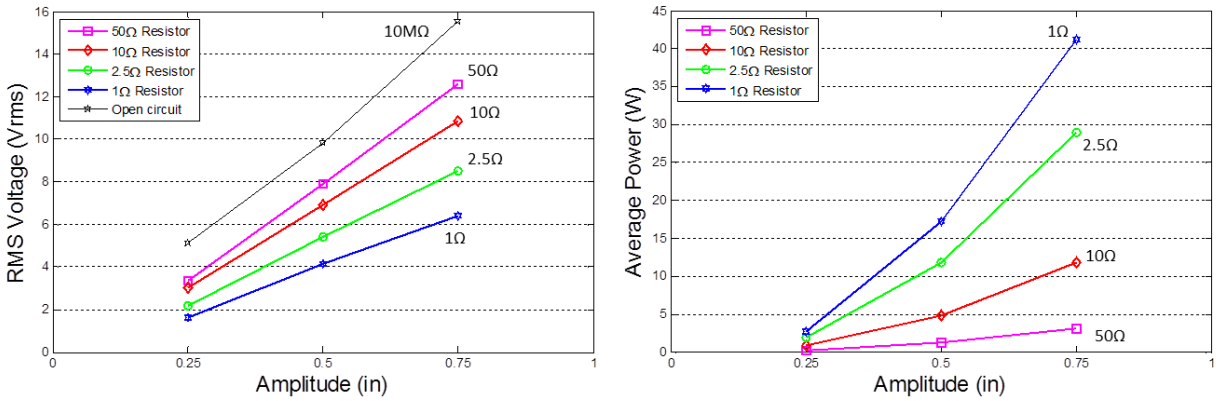

Figure 4.14: First prototype with clutch bearings - Tests at 1Hz

It is important to mention that part of those improvements can be attributed to the different gear ratio. As mentioned before, the system has a different gear ratio between compression (25:1) and rebound (20:1). Before the implementation of the clutch bearings, it was a constant 20:1. This means that now in compression, the generator spins 25% faster. This change alone can then be credited for about a 12.5% increase in voltage.

In parallel with the spring and the damper, the system will impact the behavior of the suspension. Beyond unavoidable friction, the two main forces generated are functions of the speed and of the acceleration. Whereas the voltage generated is proportional to the speed of rotation of the generator, the current is proportional to the torque. Since the load is purely resistive, the current is also proportional to the voltage, according to Ohm's law, leading to a theoretical linear relationship between force and velocity. The inertia of the different parts will lead to larger forces in the acceleration phase and reduced forces during deceleration, similar to what is experienced with an inerter. The actual force vs. velocity from tests with a 10Ω resistor (Figure 4.15) clearly shows that behavior.

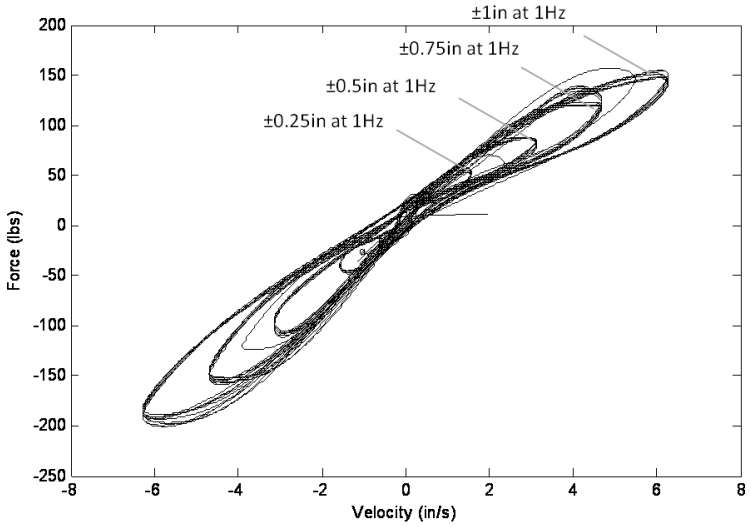

Figure 4.15: Experimental force/velocity results

This design is definitively more sensitive to accelerations than the linear generator concept. If they are high enough, the accelerations could lead to the ball screw buckling. Rated for 750lbs in dynamic loading (4,000lbs for static loading), the screw has never experienced more than half of that limit during any of the tests. The fact that the generator always rotates in the same direction also helps to reduce the forces and limit the impact of a sudden acceleration, compared to a system where the generator is directly linked to the ball nut.

4.3.2.4 *Conclusion*

The implementation of the clutch bearings and the elegant way in which the planetary gear causes one-way rotation of the generator lead to good improvements both in terms of electrical output and mechanical input. The output voltages remain on the low side, and the small generator requires more speed than it is able to receive. The next step is to improve the heart of the mechanical-to-electrical transformation, the generator.

4.3.3 Second Upgrade: Custom Generator

4.3.3.1 *Concept*

So far, an off-the-shelf RC airplane brushless DC motor has been used; its Kv of 740 rpm/V is still relatively high. A new improvement to the prototype is to modify the generator to the specific needs of the project. The goal is to lower the Kv value of the generator to be able to increase the output voltage for the same inputs, while at the same time removing one stage of the gearbox. Up to that point, the planetary gearbox was composed of two distinct stages, and removing one of them should decrease the mechanical losses and reduce the inertia, thus increasing the system efficiency. The gear ratio would be lowered to 5:1 in compression and 4:1 in rebound.

4.3.3.2 *Custom Motor*

A different motor is used as a base for the new generator: the S-3026-8 from Scorpion Power System. This motor is slightly larger in diameter (1.45in) and longer, but it still fits, without any modification, at the exact same location in the prototype, thus maximizing the space. From the factory, the motor has a Kv value of 1212 rpm/V. The exact same motor with different wiring is also available with Kv's of 980 and 840 rpm/V. All these values are higher than the previous motor, but since the idea is to completely remove the wiring to redo it, it is not an important factor in the choice of what motor to use. Its construction, its general dimensions, and the size of the magnets are factors that led to choosing that particular model.

The idea is to decrease the Kv by at least a factor of 5 to compensate for the lower gear ratio. To lower the Kv, the new wiring needs to include a larger number of turns per coil. As explained for the linear generator, the larger the windings, the higher the voltage. The choice is then to wind the coils with 30 turns of 28 AWG wire (up from the original 8) on a DLRK pattern and to connect the three phases in a wye configuration instead of a delta (Figure 4.16).

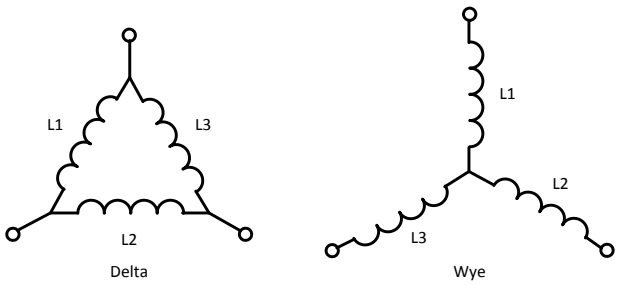

Figure 4.16: Delta and wye configuration

Connected to a central point (wye pattern), two phases, then in series, contribute to the output voltage instead of only one when connected in delta. The phases are 120° apart so the amplitude of the sum of two phases in series is not directly the sum of the amplitudes of each, but rather it increases by a factor of $2 \cdot \cos(30°) = \sqrt{3} = 1.732$.

Similarly to the technique presented in the previous chapter, the output voltage and thus the Kv value can be estimated through a magnetic analysis using the software FEMM. The model of the motor has already been created in the 3D CAD software SolidWorks and can be imported in FEMM, simplifying and accelerating the creation of the magnetic model (Figure 4.17).

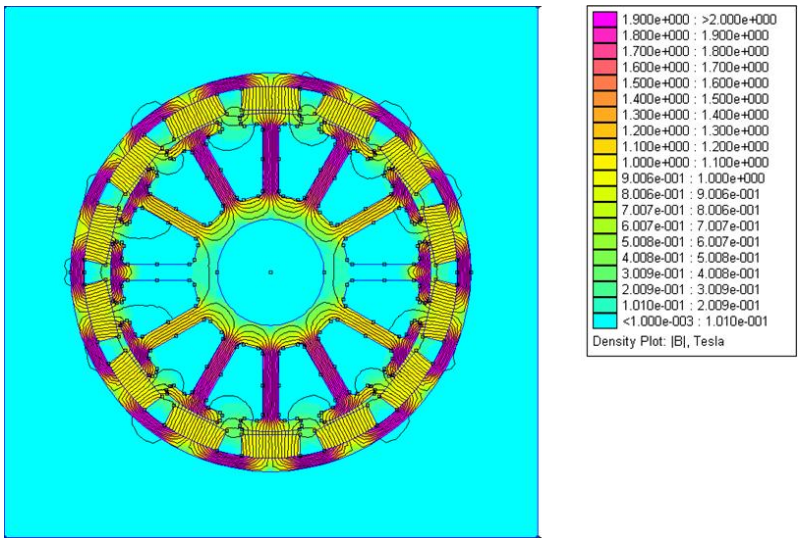

Figure 4.17: Finite-element magnetic analysis of the generator

The model is then analyzed in various angular positions of the rotor, and the flux for each pole is measured and recorded (Figure 4.18). Based on the chosen wiring pattern (DLRK), the flux is summed for each phase, and using Faraday's law, the estimated voltage for each phase is computed. With a wye

configuration, the output voltage is the sum of two phases. The rectified voltage after the diode is also presented.

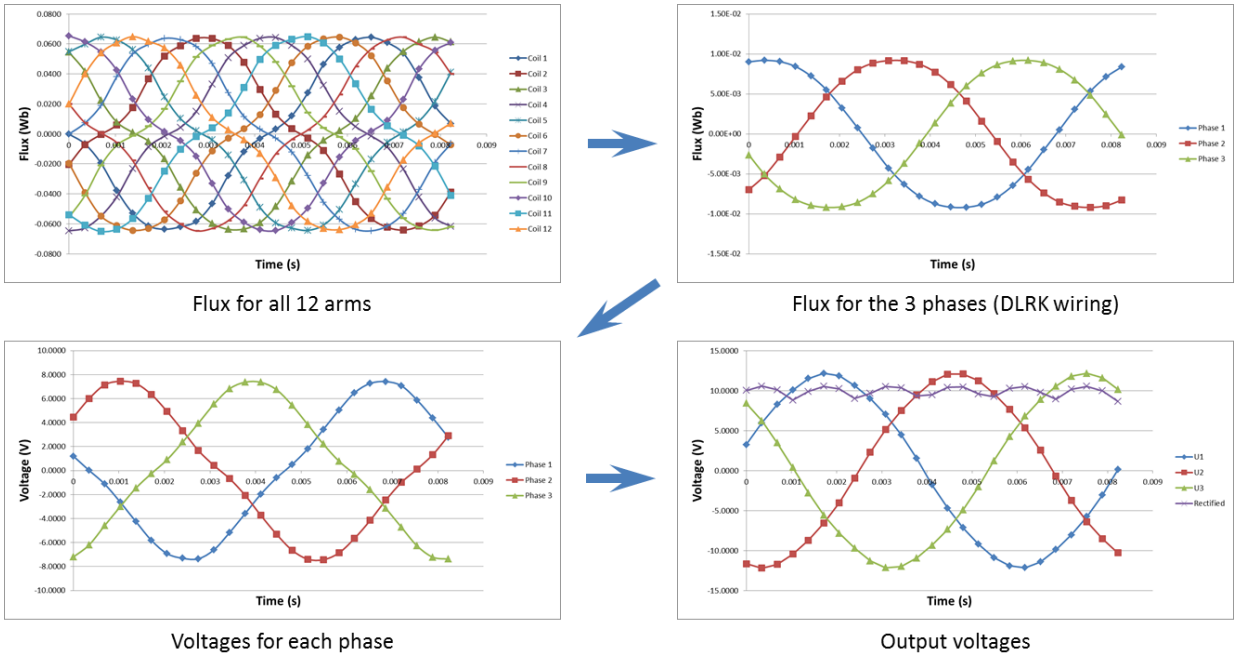

Figure 4.18: Post-process of the magnetic analysis results to determine the generator Kv -1000RPM

The analysis shows that the customized generator has an estimated Kv value around 90 rpm/volt. This is then confirmed by tests on the actual generator (Figure 4.19). This means that the Kv is about 13 times lower than from factory (1212rpm/volt), and about one-eighth of the previous generator used.

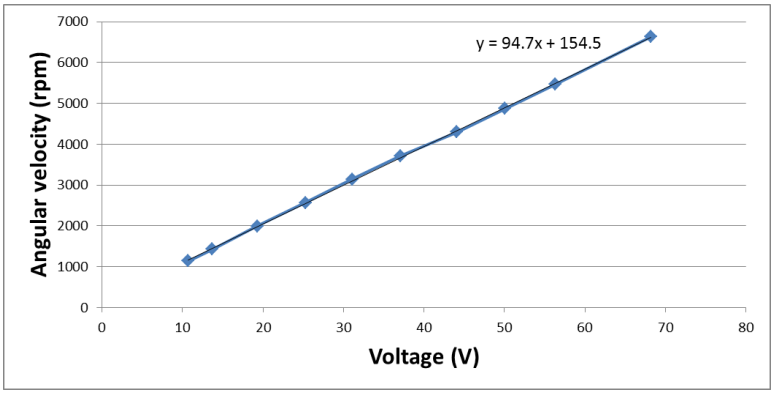

Figure 4.19: Open-circuit voltage characteristics

A rig was specially built to test the generators (Figure 4.20). A brushed DC motor was used as the prime mover, driving the generator under test through a flexible coupling. The generator is not directly attached to the chassis, it is free to rotate. An arm attached to the stator stops the rotation by pushing onto a load

cell. It is then possible to compute the torque necessary to hold the generator still, which, at a constant speed, is equal to the input torque supplied by the DC motor, neglecting the friction in the bearings. This is an inexpensive alternative to an actual torque transducer. A laser tachymeter is used to measure the speed of rotation. With these two mechanical parameters, torque and speed, the input power can be calculated. The output power is, as previously, obtained by measuring the voltage across a resistor of known value. With the power in and out, it is possible to compute the efficiency of the generator.

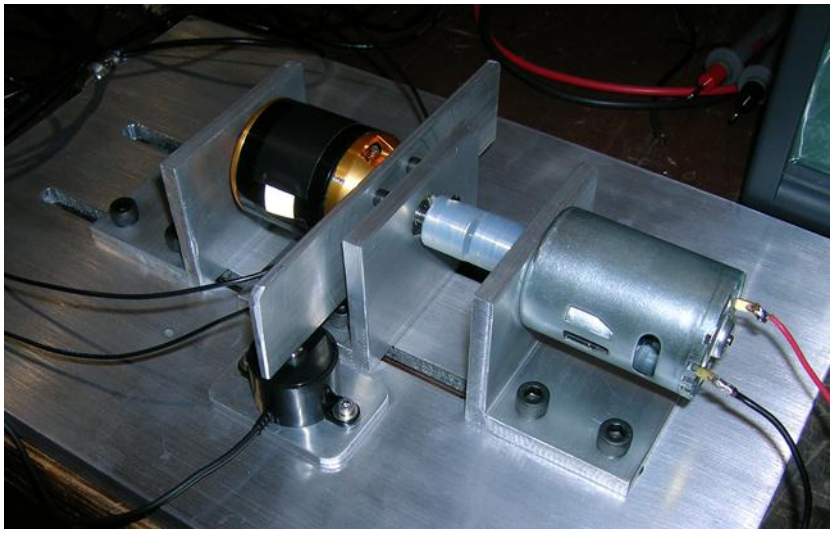

Figure 4.20: Generator test rig

One issue is that the rig is underpowered. With the new wiring, the motor is able to generate a decent resistive torque even at low speed. Therefore, low resistance could not be tested. However, with a 50Ω load, the generator is already able to produce 40 Watts at 5000 rpm, with an efficiency of about 75%. There is no doubt that the generator is able to reach higher efficiency, as brushless motors are known to have efficiencies in the mid 90%.

It must be reiterated that the drawback of a larger number of turns in the coils is the increase in the internal resistance, but for the range of voltages and currents expected, it remains totally acceptable, with a value of about 1.7Ω.

4.3.3.3 *Prototype*

The main change in this upgrade is obviously the generator. A few parts had to be modified to fit around it, such as a stator mount and a bearing support. A special shaft also had to be machined. But besides those few elements, the rest of the prototype stays exactly the same: identical ball screw and the same clutch bearing mechanism to make the generator rotate in only one direction. The changes are really limited to the customization of the generator, several new mounting parts, and the disappearance of one of the two stages of the gearbox.

The generator sub-assembly (Figure 4.21) was then slid in place and the energy harvester was ready for a new series of tests.

Figure 4.21: New generator sub-assembly

4.3.3.4 *Prototype testing*

The prototype was then tested again on the shock dynamometer, in the same conditions: all tests were conducted at 1Hz. The amplitudes range from 0.25 inch to 1 inch. The loads are open circuit, 50Ω, 10Ω, 5Ω, and 2.5Ω. Lower resistances were not tested to avoid reaching currents that could damage either the generator windings or the electronics. The current was already reaching $5A_{RMS}$ and $7.7A_{peak}$ with the 2.5Ω load.

The open-circuit voltage (Figure 4.22) now reaches 20 V_{RMS} for a 0.75-inch amplitude at 1Hz, and over 6V_{RMS} for the smallest input. Improvements are also seen for the voltages across a resistive load, with about 30% and 15% higher voltages for a 50Ω and a 10Ω load, respectively. The average output power (Figure 4.22) evidently follows the same trend: an increase of 70% with the 50Ω resistor, and 30% with the 10Ω load. This shows good improvement over the previous generator, especially considering that the generator is now turning five times slower.

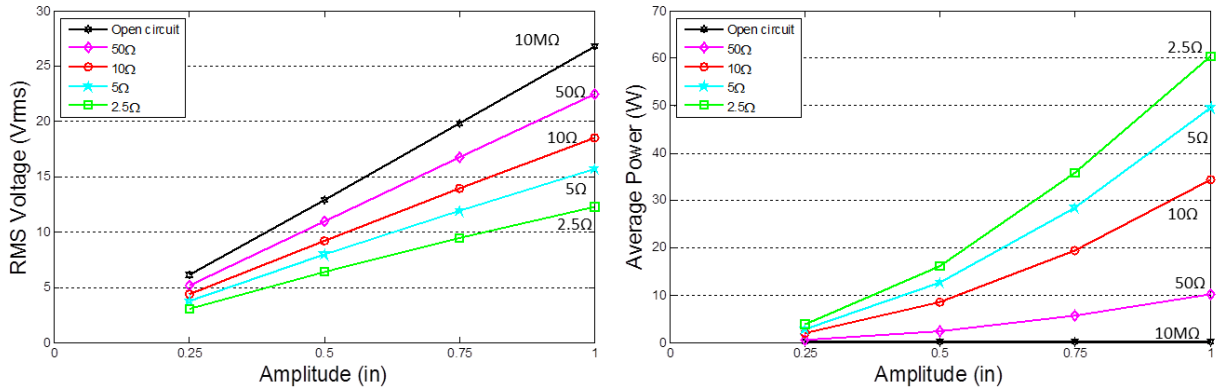

Figure 4.22: First prototype with clutch bearings and custom generator - Tests at 1Hz

All the values show improvement, although the raw absolute maximum power is lower due to the lack of tests with the 1Ω resistor, in order to protect the system from unreasonably high current. A test with an input of a 1-inch amplitude at 1.5Hz and a 10Ω resistor leads to an average output power of 65Watts, proving that higher power can definitively be achieved.

4.3.3.5 *Conclusion*

The voltage and power increases are from the previous design, in addition to the improvements resulting from the implementation of the clutch bearings. So, compared to the very first rotational harvester, the performances are significantly better. There is no reasonable comparison to be made with the output power of the linear generator.

The concept is proven to be very promising, and it will be further developed to exactly match the need of the railroads.

4.3.4 Double Generator

In noticing that there was some space available in the system, the idea developed that a second generator could be included, driven by the same shaft as the generator already in place. This would allow more power to be produced without significantly increasing the mechanical losses, thus improving overall efficiency.

The modifications to the current prototype were quite limited: only a new shaft and stator mounts were required. The second generator was subjected to the same rewiring process. But this time, two strands of wire were placed instead of only one (Figure 4.23). This does not affect the Kv value, but it divides by two the internal resistance, thus reducing the ohmic losses by 50%. This results in lower decreases of the output voltage due to low resistance and higher current capabilities.

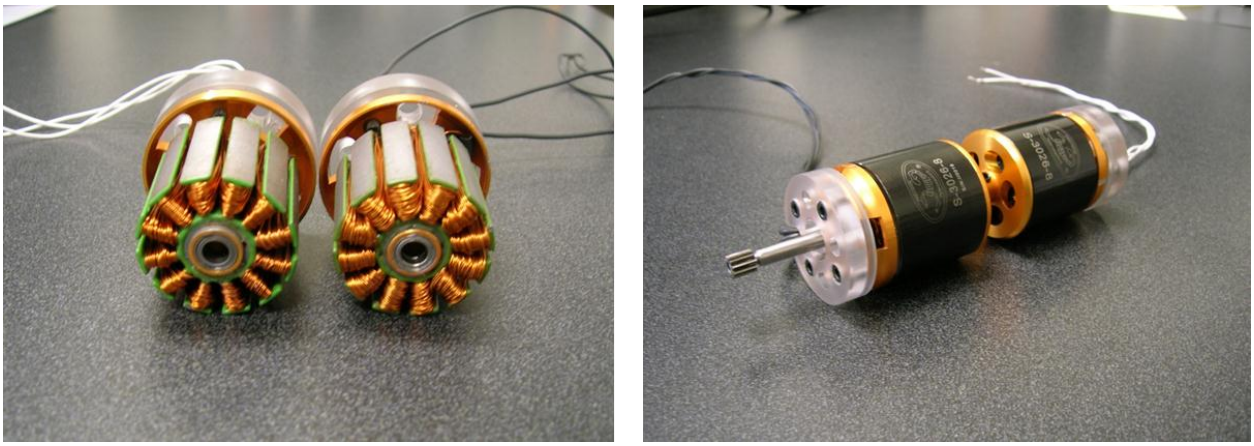

Figure 4.23: Generators with different wiring (left) and new generator assembly (right)

The system has not been fully tested on the shock dynamometer, but qualitatively, in the table demonstrator (Figure 4.24), the system is now able to reach higher voltages. The force to move the system is also significantly larger. The two generators have their own diode bridge, whose outputs are placed in series. There may be a disruptive effect of one generator on the other's diode bridge, leading to potential intermittent short-circuits. The two generators may need to be synchronized to fully benefit from having two of them.

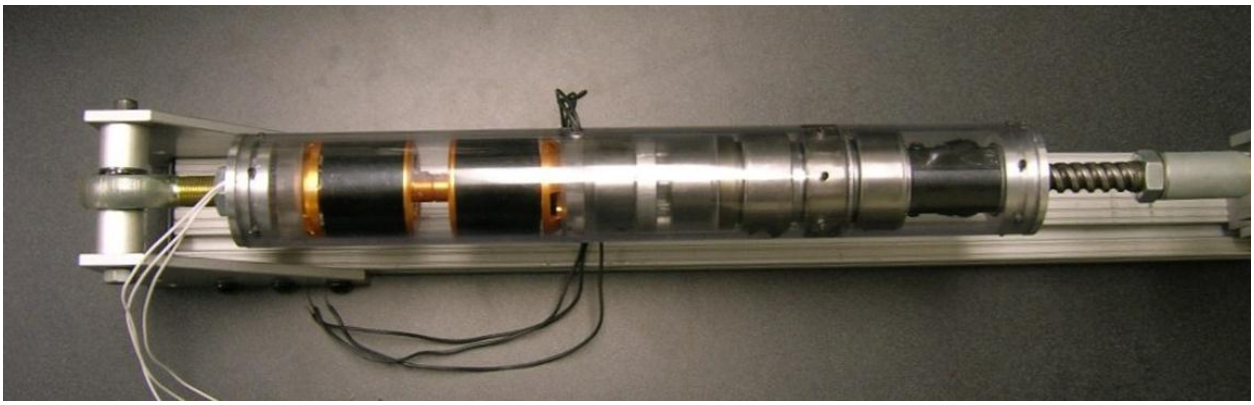

Figure 4.24: Energy harvester using two generators

4.3.5 Conclusion

From the very first prototype, the use of a rotating generator proved that it offers great potential and that the possibility of having a gearbox to increase the speed is key for good performance, leaving the concept of a linear generator far behind. Two successive improvements were made to the prototype and led to increased performance, with larger voltage and power outputs and higher efficiencies. Having the generator always rotating in the same direction, and having a generator with a low Kv are key concepts that will be incorporated into the next generation of prototypes.

4.4 Second Prototype (for D5 Springs)

4.4.1 Introduction

Although the energy harvester is similar in size and shape to a classic coil-over shock absorber, a freight car suspension is no ordinary suspension: the spring is significantly different from an automotive suspension spring. According to AAR's *Manual of Standards and Recommended Practices*, a D5 truck spring (Figure 4.25), which is the standard for many American freight cars, has a large inside diameter at 3-19/32 inches for a relatively short free length with "only" 10.25 inches. The deflection from free to solid length is 3-11/16 inches, which yields a compressed length of 6-9/16 inches. At that point, the spring is under a force of 8,266 lbs. Typically, eight or nine of those springs are placed at each corner of the car. To increase the spring rate, inner springs can also be used. Although there are different places with linear motion on a freight car, the main goal has always been to use the motion of the suspension and the free space inside the spring. But as it is, the current prototype, with a length of about 12 inches for a stroke of 3 inches, does meet the requirements. A second system was then designed.

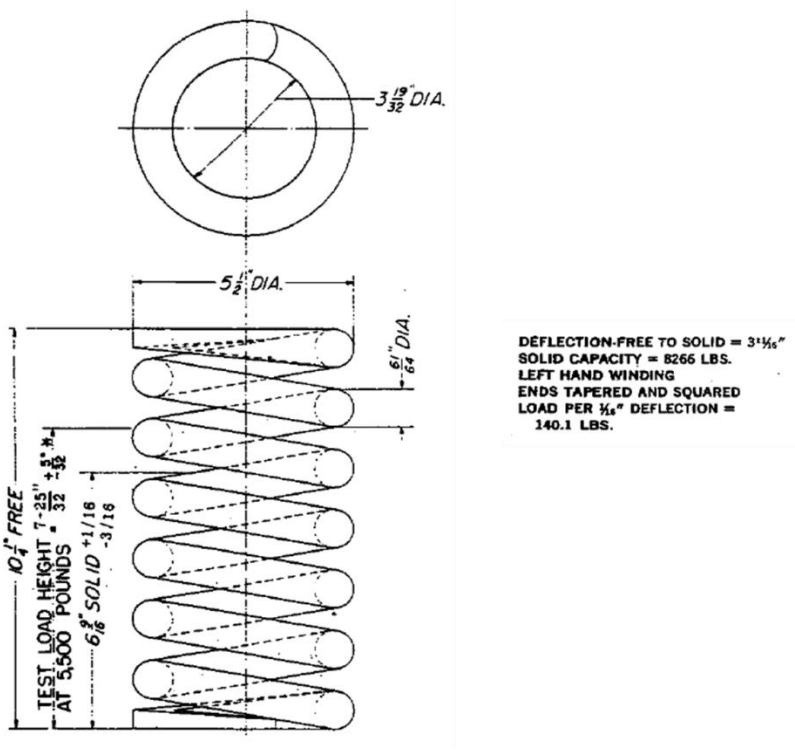

Figure 4.25: D5 spring dimensions, according to AAR's *Manual of Standards*

4.4.2 Design

Following the good results of the first generation, the new prototype ultimately keeps the same concept: a ball screw transforms the initial linear motion from the suspension into a rotation. Then an arrangement of clutch bearings and a planetary gear increases and rectifies the motion of the ball nut before finally driving a permanent magnet generator. However, although the concept is identical, there is such a difference in the dimensions that the prototype must come from a complete redesign.

With a required stroke of almost 4 inches and a maximum length of 6.5 inches, and considering that the ball nut is 1.5 inches long, the ball screw must be able to pass through the planetary gear and the generator. In the previous prototype, those elements are too small and it is physically impossible for a half-inch rod to pass through them. The D5 spring greatly reduces the available space lengthwise, but it allows a system with a diameter twice as big. So, all the parts need to be shorter, but can be thicker. The generator used for the new prototype is a MonsterPower 160 (Figure 4.26) from the same manufacturer as the very first generator: Exceed RC. It is also a permanent magnet outrunnner brushless DC motor, designed for RC airplanes. Its shaft is 10mm (0.4 inch) in diameter, which is a significant increase in size compared to the generator previously used, but that still makes it too small to directly fit the ball screw. However, it is possible to modify the motor, using different bearings to gain enough room.

Figure 4.26: Exceed RC Alpha 400 (left) and MonsterPower 160 (right)

The ball screw remains identical to the previous prototype: 0.5-inch diameter, 0.5-inch pitch, an acceptable dynamic loading of at least 750 lbs (up to 1,300 lbs from some sources), and a static rating at 4,000 lbs. In all the tests conducted on the previous prototype, the forces never came close to those limits, and there is confidence that the ball screw will also perform well in the new system. Care is also taken in

the design so that the ball screw is sliding in one single part (beside the ball nut), eliminating junctions where the end of the screw could hit and be jammed.

All the elements of the system, i.e. bearings (one-way and regular), and planetary gears, experience a significant growth in size. Densely packed, they manage to fit inside a 6.5-inch long, 3-inch diameter tube (Figure 4.27).

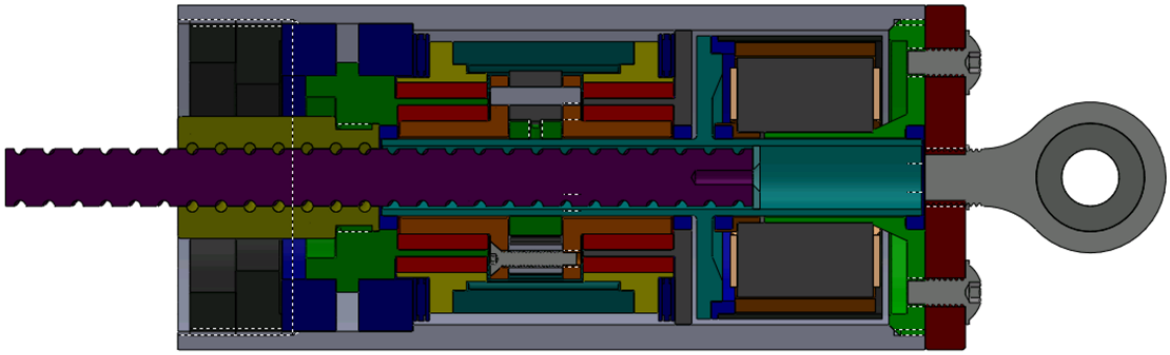

Figure 4.27: Final design of the prototype of second generation

With a 5-inch stroke, the design meets all the requirements set by the dimensions of the spring. It can accommodate the extended length of the spring and still be short enough to be protected by the spring in full compression (Figure 4.28).

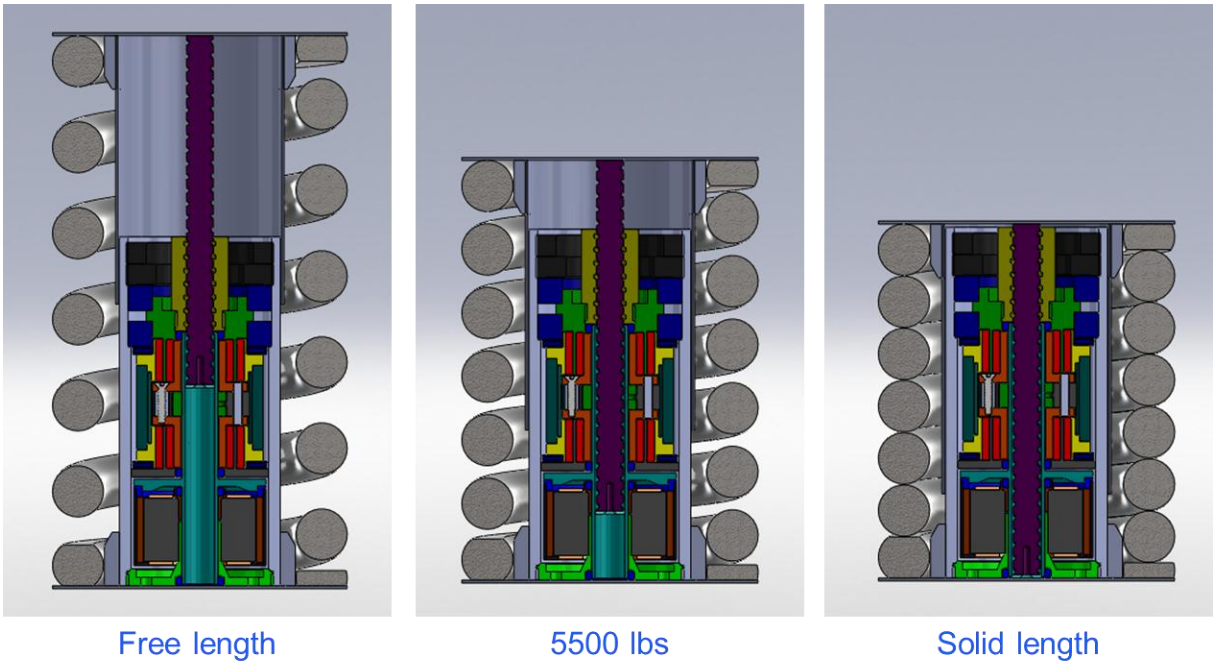

Figure 4.28: System placed inside a spring fully extended, under test load, and fully compressed

The importance of a generator adapted to the particular application has been demonstrated in the previous section. So, for the new prototype, the generator was also customized with completely new wiring. The generator has a Kv from the factory of 245rpm/volt. This is fairly low, but the application needs an even lower number, especially since the planetary gearbox now has a ratio of 2:1 or 3:1, depending on the direction of rotation. A finite-element analysis of the magnetic characteristics of the generator was then performed (Figure 4.29), and it resulted in a value of 1377rpm/volt/turn for the Kv value. This led to the decision to have 35 turns per coil in order to obtain a Kv of about 39 rpm/volt that would yield about 12 volts at peak for a ±0.1875-inch amplitude at 1Hz input.

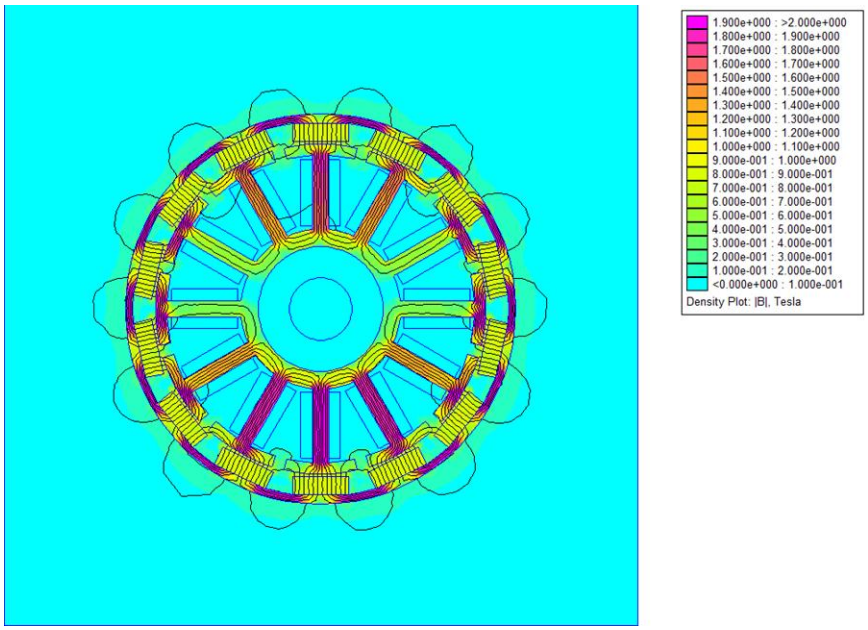

Figure 4.29: Finite-element analysis of the generator magnetic characteristics

It is worth noting that the field lines and the greenish ring around the generator on the simulation are proof that the magnet ring is undersized and should be made thicker. This somewhat decreases the efficiency of the generator and can lead to eddy currents in the nearby outer tube.

4.4.3 Prototype

The ball screw/nut, the planetary gears, and the generator are, as for the previous prototype, off-the-shelf components. A dozen parts were custom made for the new prototype, and outsourced to a professional machine shop. With all the parts received, modified, and prepared, two prototypes were completely

assembled. The dimensions sharply contrast with the previous prototype (Figure 4.30): the larger diameter, up to 3 inches from an initial 1.5 inch, makes an impressive difference.

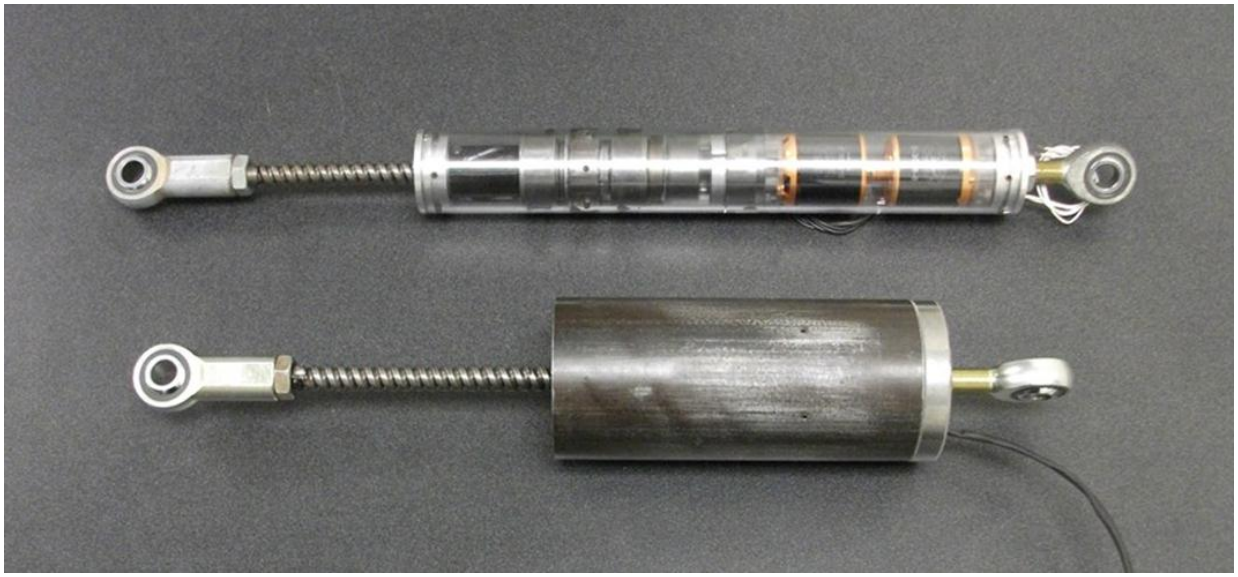

Figure 4.30: First (top) and second (bottom) prototype generations

Early tests showed that few modifications were needed. First, the sun gear was glued in place and also held by a set screw. But under heavy repeated loads, it managed to come loose. A flat on the shaft corrected this issue. Similarly, the part holding the clutch bearings that prevents a backward rotation of the planet carrier and the ring gear was maintained in place with set screws, but was not enough to keep it from sliding. Notches were made for the set screws to rest in, and the sliding issue was solved.

A third issue was that large forces were observed at certain positions. This was traced back to misalignment of the thrust bearings due to the bearings themselves or to the machining of the bearing seats. A slightly compliant element, a rubber washer, was then placed between the bearing and the tube, greatly helping to soften the hard spots.

4.4.4 Characterization Tests

With the prototypes fully assembled, they were tested, in the same conditions as the previous prototypes, in the laboratory's shock dynamometer (Figure 4.31). Although the test rig is capable of producing in any type of input, the characterization tests were again limited to purely sinusoidal motions. They are largely

sufficient to characterize the performance of the system, and are of great value to compare the different prototypes and quantitatively measure the evolution of the successive prototypes.

Figure 4.31: Second generation prototype in the shock dynamometer

The tests again primarily focused on those conducted at 1Hz. Amplitudes of 0.25, 0.5, and 0.75 inch were once more used, but 0.1 and 0.1875 inch now completed the list. Three resistors, beside the open circuit, were used to draw different levels of current.

The new prototypes show great improvements both in terms of voltage and power (Figure 4.32). Now, it is possible to obtain an open-circuit voltage over 12 V_{RMS} with only ±0.25-inch at 1Hz. Across a 5Ω resistor, this voltage drops slightly, to about 7.1V_{RMS}, which brings over 10 Watts of power (average). This is the range of values that had been desired since the beginning of the project. With larger inputs come larger outputs. With a 0.75 amplitude, the system generates voltages between 20 and 37V_{RMS} (5Ω and 10MΩ, respectively), and average output powers up to 85 Watts. The peak power is recorded at about 230 Watts.

Low amplitude tests have also been conducted. The prototype shows that a displacement as little as 0.1 inch can lead to almost 6 volts generated in open circuit, or 2.75V_{RMS} with a 5Ω resistor, leading to about 1.5 Watts produced on average. This is quite good for so little motion.

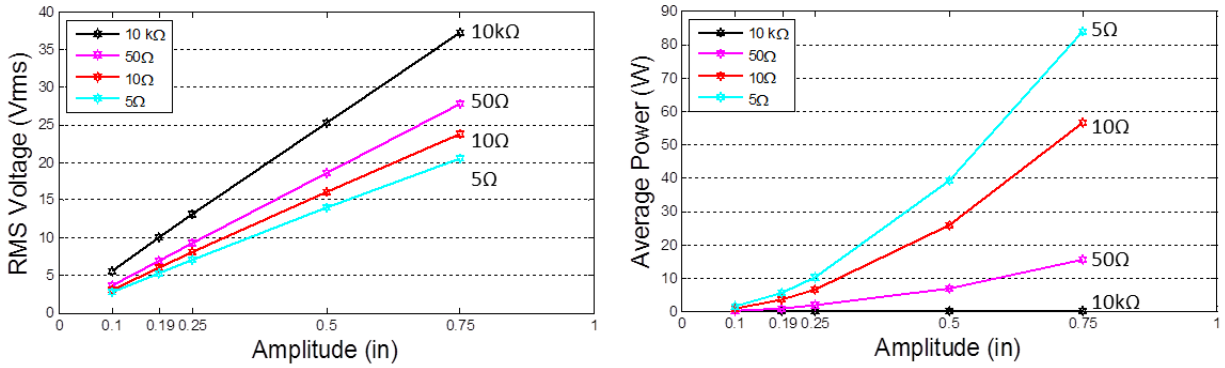

Figure 4.32: Characterization tests at 1Hz

Because of the linear relationship between the speed and voltage, it should be possible to obtain the exact same range of voltages and powers with only half the amplitudes if the frequencies are doubled. So, tests were conducted at 2Hz, with amplitudes of 0.1, 0.1875, 0.25, and 0.375 inch (Figure 4.33). The curves at 1Hz and at 2Hz show an excellent match. For example, with the 50Ω resistor, the RMS voltage reaches 27.85V_{RMS} with the 0.75inch amplitude (1Hz), and 28.45V_{RMS} with 0.375 inch at 2Hz: a 2% difference. Also, the maximum average power measured with a 10Ω resistor is 56.42 Watts with 1Hz inputs, and 56.73 with 2Hz: a 0.5% difference.

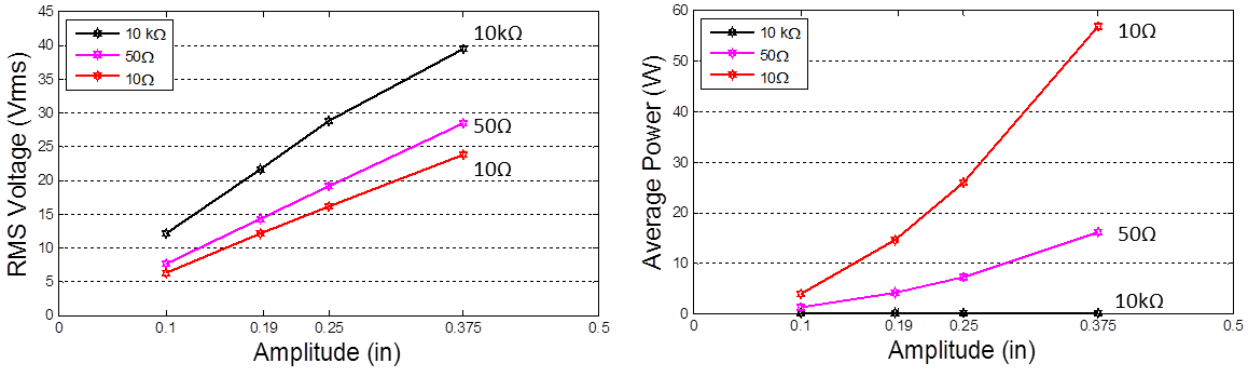

Figure 4.33: Characterization tests at 2Hz

All the test results were highly repeatable. The outputs recorded one day could match the next day or months apart. Furthermore, both prototypes produced amounts of power that are quite close, generally within one percent of one another.

4.4.5 Durability Tests

The tests to characterize the performance of a prototype are usually relatively short: generally 10 cycles for each amplitude. So, the next step was to run some durability tests to verify that the performance of the energy harvester can be maintained over long periods of time. The ultimate goal is to subject the prototype to over two million cycles that would represent a few years of service. It was planned to use the crank dynamometer to perform those tests. Whereas the eletro-magnetically actuated (EMA) shock dynamometer uses a linear motor to provide the motion, the crank dynamometer uses a more conventional motor. The rotation is then transformed into linear motion by a scotch yoke. With the steady input that is foreseen for the durability tests, this is a more appropriate solution. In the installed software for the test rigs, there was no dedicated function for durability testing, so functions dedicated to other tasks, such as the warm-up sequence, were repurposed to generate the input. The very first test for durability unfortunately ended very quickly. After only 7 minutes, the system completely locked up, managing to move the top cross-bar. The failure was tracked down to a seized planet gear. Whether the gear had a defect or the pin was too soft, the result was a large tear in the pin and the gear completely locked on it, preventing any motion from the entire system. Stainless steel was, a posteriori, not such a good choice for the material for the planet gear pins, then replaced by hardened steel pins. The other gears were checked for defects but none were found.

Figure 4.34: First failure - seized gear

Although the second prototype lasted 2 hours without any issues, the need for a safety feature to stop the test if the force exceeds a given limit was now clear. Roehrig Engineering offers a software solution for durability testing on any of their test rigs but that was not, for us, an economically viable option. The idea was then to move the tests to a hydraulically actuated MTS material test rig available in the CVeSS

laboratory (Figure 4.35). Dedicated to traction and compression tests, it is not a shock dynamometer but it provides the same possibility of replicating the linear motion of a suspension. The actuator is controlled by a MTS analog controller (model 458.20) which is capable of producing a large variety of inputs including square, ramp, and sinusoidal waves of any amplitude and frequency.

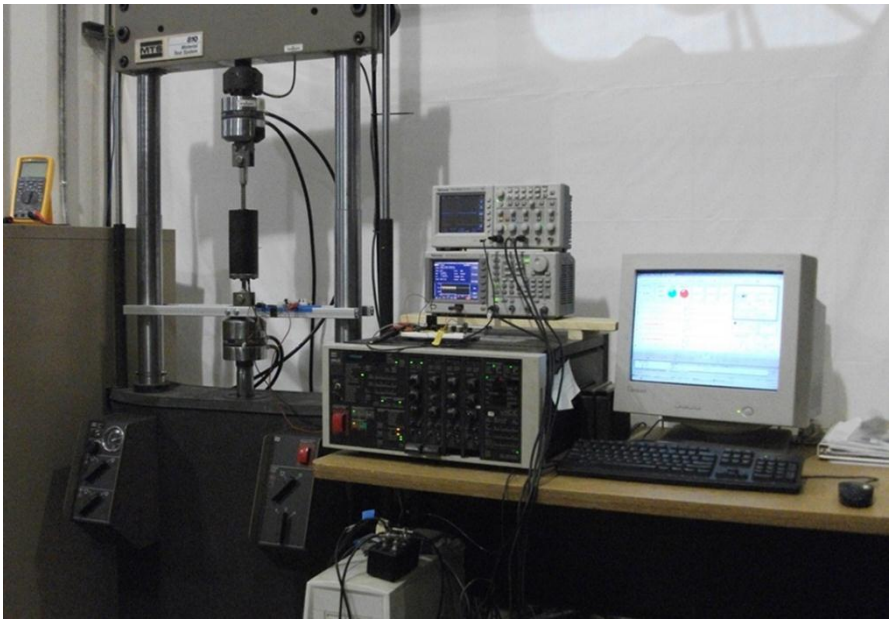

Figure 4.35: Test setup with the MTS test machine

Controller command, actual displacement, and force are measured from outputs on the controller with a PC-based dSpace data acquisition system. Along with the mechanical conditions of the test, a fourth input is ultimately used to record, through a voltage divider, the voltage across the resistive load connected to the energy harvester (Figure 4.36).

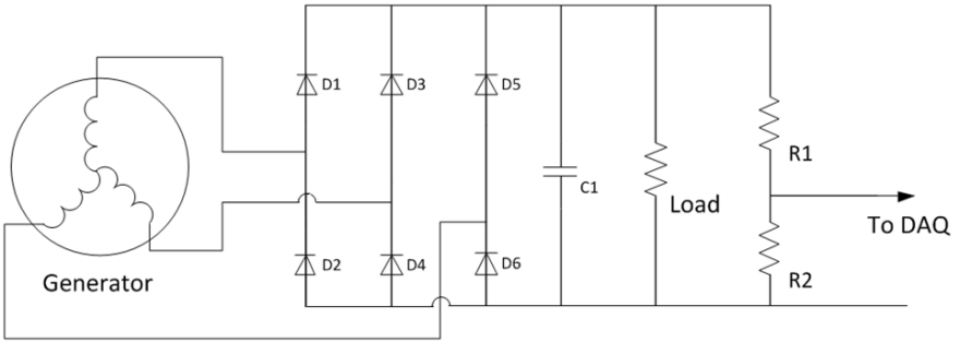

Figure 4.36: Electrical schematics used to record the generator output

The dSpace system works in conjunction with Matlab/Simulink and Control Desk software. Simulink gives the users the option to define what data to collect and after what treatment. Corrective factors can be applied to account for the scales of the MTS controller or the voltage divider ratio. The sampling frequency is also defined here. Almost all tests were sampled at 1000Hz, later downsampled in Control Desk to 500Hz. It is also possible to instantly derive other test parameters like the output current and power. Simulink also allows more complex functions that are required to conduct a durability test. First, it would be impractical and useless to record data continuously through the 2 million cycles. Instead, ten seconds are recorded at regular intervals, usually every 2 minutes. Having a clock that starts with the simulation is also a necessary function. The internal clock of dSpace starts right at the time the model is built. A second clock is then incorporated to start with the test, and every recording is based off the time it gives. If the force limit were to be exceeded, there is a function that would automatically trigger data recording of the event to have more insight into what went wrong.

Control Desk provides a user-definable graphic interface. It allows display of any variables present in the Simulink model; change of some of the parameters, like the value of the resistive load for instance; and it also saves the data. The layout is highly customized (Figure 4.37) to have every function needed during the durability test. It displays the time since the start of the test and since the model was built. This is important information because the dSpace clock stops after 4 hours and 40 minutes. The model then needs to be rebuilt.

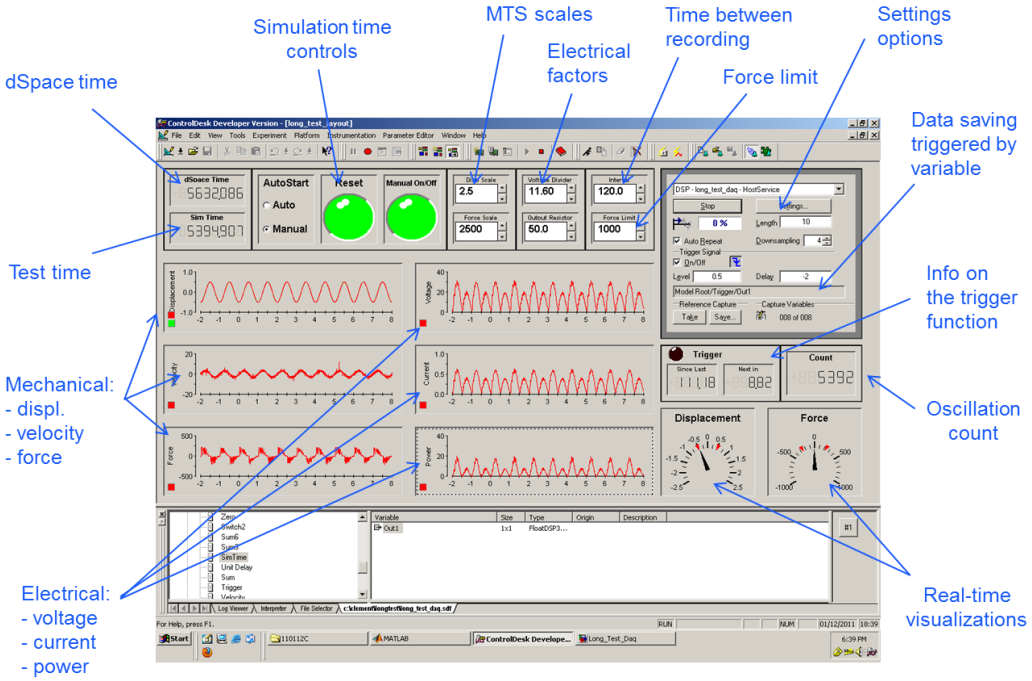

Figure 4.37: Custom dSpace layout

It also displays the time plots from the last 10-second sample of all the mechanical conditions (displacement, velocity, and force) and the electrical outputs: voltage, directly measured by the DAQ board, and the current and power derived in Simulink. Two clocks show the time since the last sample and when the next one will take place. A counter keeps track of the number of oscillations for the running tests. The time plots are only refreshed when data is recorded every two minutes, so to have an idea of the displacement and the force in real time, two dials are present in the layout. The important factors of the test, like the MTS scale factor, force limit, and resistive load, are also shown. They can directly be modified from here without having to rebuild the Simulink model. The last element present is the settings options panel, where the user can define what variables are saved: the triggering function parameters or the downsampling ratio.

With the Simulink model coupled with the custom Control Desk layout, the durability tests are ready to be fully recorded. But before starting the test campaign, one last thing is required. Since the system uses a ball screw, when force is applied, a reactive torque is generated. This is not an issue with the Roehrig shock dynamometer, as the linear motors are guided and prevent any rotation of the shaft. However, the MTS test rig is actuated with a hydraulic cylinder, which means that a rotation becomes possible. To prevent this from happening, a special bracket is made (Figure 4.38). Two arms extend from the mounts to slide against the vertical posts of the rig, and are protected with plastic sleeves. The design uses two arms so there is virtually no radial force on the hydraulic cylinder.

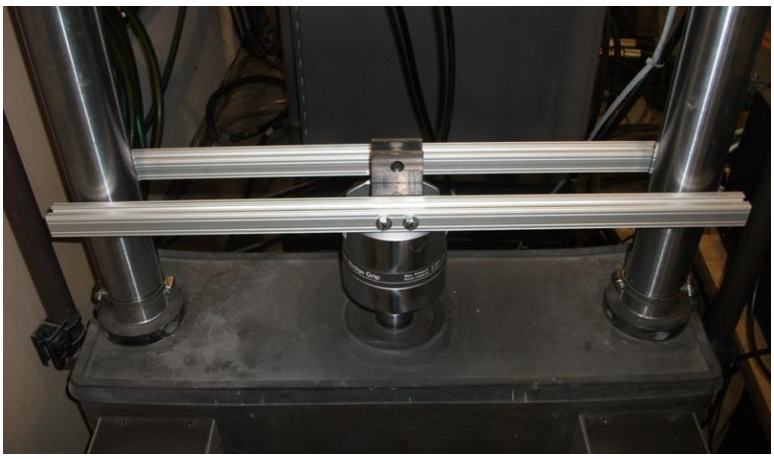

Figure 4.38: Custom bracket to mount the harvester in the MTS test rig

The same input is used during the entire durability test campaign: an amplitude of 0.5 inch, i.e. 1 inch peak to peak, and a frequency of 1Hz. This is somewhat in the middle of the range of all the tests conducted so far, and it makes the results easier to scale. Two different loads are used: a 50Ω resistor and

a 20Ω resistor. After a short period to fine-tune the testing procedure, a typical test would last 12.5 hours, leading to 45,000 cycles per day, usually five days per week. Every two minutes, ten seconds of data were recorded. The voltage and power can be averaged to yield just one representative number for each parameter, for each sample. The values for the RMS voltage and average power of the nearly 16,000 samples recorded during the three months of testing can be seen on the two plots below (Figure 4.39).

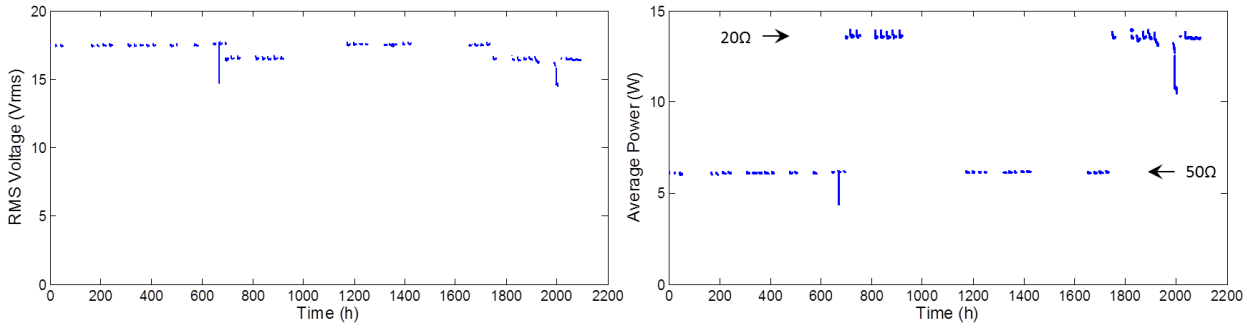

Figure 4.39: Average voltage and power over the entire durability test campaign

It may appear unclear, but the line on these two plots is solid for the whole length of a test. When the test stops (such as at the end of the day), the line also stops. So what may appear as a dash or a small group of points actually represents up to about 500 points, usually for an entire day of data. Starting with the very first test and ending with the very last, the time scale is 2200 hours long, which is equivalent to about 3 months. The plots show the entire time period to put the data in its context, with variable gaps between tests.

The output voltage and the input force are rather constant throughout a test day. Both have a tendency to decrease during the first hour of test before becoming quite stable. The voltage only drops by less than 1%, whereas the average force decreases by about 15%, ultimately raising the efficiency. As can easily be seen on the plots, there are two obvious ranges of values. They are related to the two different loads, 50Ω and 20Ω, connected to the generator at different points of the 2-million-cycle test. With the 50Ω, the voltage remained around 17.5V_{RMS} and the average power at 6.1W. With the 20Ω resistor, those numbers become 16.5V_{RMS} and 13.5W, respectively. With a 50Ω load, the efficiency is consistently around 34% and climbs to about 50% with the 20Ω resistor.

The 2-million mark was reached in just under three months (Figure 4.40), with the prototype still functioning properly and without any major problems. Two million cycles at 1Hz corresponds to 555.5 hours of tests, which equals 23.1 days.

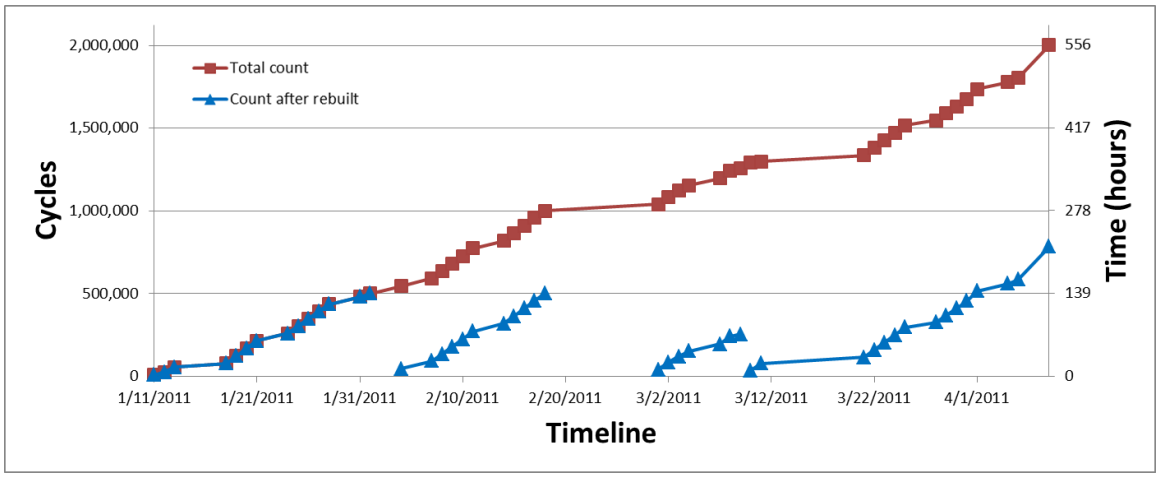

Figure 4.40: Increase of the cycle count over time

At three different times during the two million cycles, the prototype was completed taken apart and inspected for wear, cracks, or failed parts (Figure 4.41). As expected, some wear was observed, but nothing abnormal was ever detected. Furthermore, no parts (gears, ball screw, bearings, or custom-made parts) failed or needed to be replaced during the course of the test campaign (or later).

Figure 4.41: Inspection of the energy harvester after two million cycles

Although no part was ever damaged, two incidents must be reported. The first one is the failure of the bond between the ring gear and the part driving it. These two parts are slightly press-fitted and glued

together. Shortly after an inspection of the harvester, the connection between the two parts broke, leaving the system harvesting energy only in one direction. The issue is attributed to contaminant or oil present on the parts in the area where the glue had been applied. The parts were cleaned up, put back together with fresh glue, and the problem never reoccurred. A similar incident occurred with the stator, which started to break free, stripping some wires. Reinsulating the wires and regluing the stator solved the issue, which never reappeared again. Those two events can be seen in the voltage and power plots (Figure 4.39) with values unusually lower. Other than those two incidents, nothing else went wrong, and the prototype performed as expected, with performance consistent through each test.

The durability test campaign ended with a 60-hour-long test, counting 216,010 cycles. For the first several hours, forces and voltages slightly decreased and then were very steady (Figure 4.42). A 20Ω resistor was used, yielding an average output power of 13.5 Watts. During the 60 hours, the efficiency was consistently above 50% (around 51.5%). Figure 4.42 clearly shows that the average force was around 100lbs, but it is also interesting to look at the maximum force that was reached. Over the 60 hours of test, the force ranged constantly between 250 and 350 lbs. This is not very high, considering the forces involved in a freight car suspension on one hand, and the ratings of the ball screw on the other.

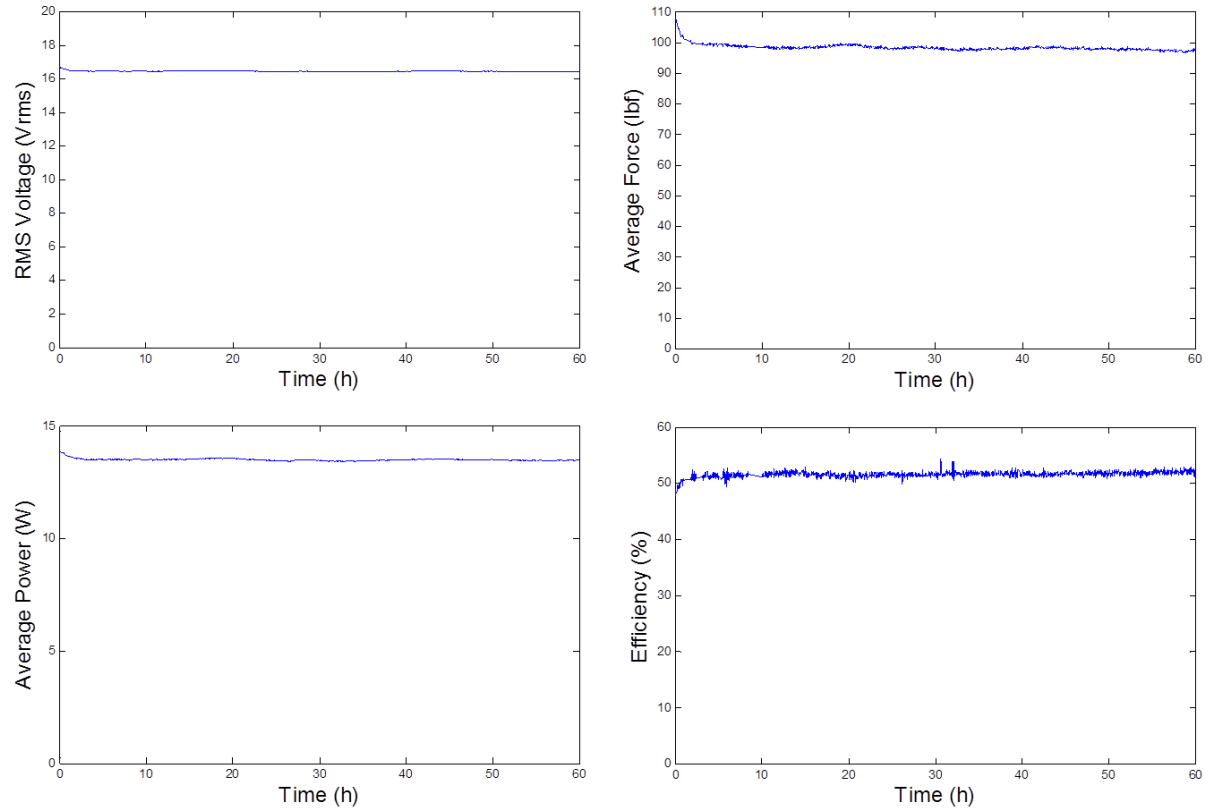

Figure 4.42: Results of the 60-hour test, concluding the durability test campaign

This very long series of tests proved that the prototype, beyond having very good performance, was also able to withstand for long periods of time the efforts involved in harvesting energy from the motion of a suspension.

4.4.6 Field Data

4.4.6.1 *Data from Amsted Rail*

Thus far, the tests have been solely done with purely sinusoidal inputs that are meant to reflect the frequencies and displacements that are encountered in the freight car suspension. The test results that will be presented in this section are obtained using displacements of the suspension provided by Amsted Rail. They are measured on an instrumented freight car while in revenue service. The set of data is actually comprised of 100 different sections of 8 seconds of data taken over the course of three days. The exact conditions of each recording are unfortunately unknown. The speed is reportedly between 20 and 60mph and the car runs in 286,000lbs gross rail load service, but its actual loading is unknown.

In order to have one coherent input motion from these 100 recordings (Figure 4.43), some modifications to the data have been made. Each section is slightly shifted up or down to match the end of the previous one, in order to remove any discontinuities that should not be present. The data set is also modified such that the end matches the beginning, allowing the test to run in a continuous manner. Some noise was present in the signal, but this was eliminated by using a low pass filter.

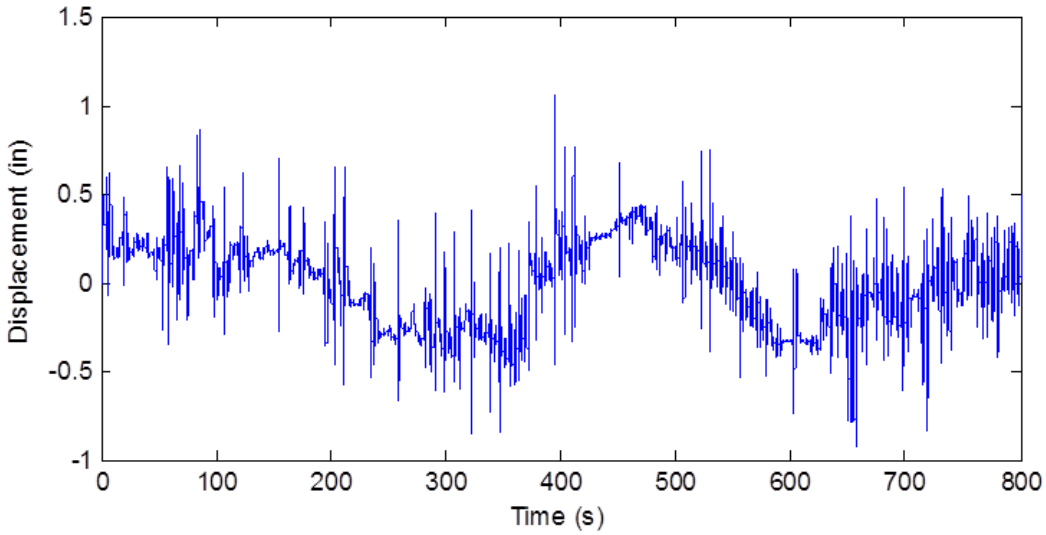

Figure 4.43: Suspension displacement from field measurements

A spectral analysis of the signal (Figure 4.44) shows that most of the energy is at low frequencies, with a peak around 2Hz, validating the range of frequencies used for the characterization tests using sinusoidal inputs.

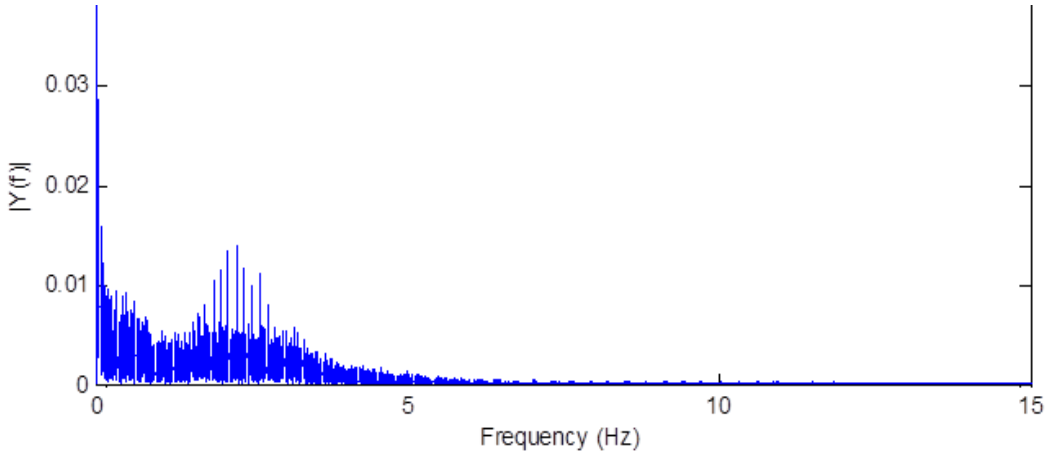

Figure 4.44: FFT of the suspension displacement

The setup with the MTS test rig coupled to a dSpace data acquisition system was once again used. The function generator of the MTS controller was not used this time because of the complexity of the input signal. Simulink does not have enough memory to store the entire input signal, so the solution was to use a function generator to provide the input to Simulink through the DAQ system. The dSpace output was then connected to the external control input of the MTS controller. dSpace was also used to provide the input profile, which allowed some extra functions like a low-pass filter to smooth the data stored externally, and a soft stop to prevent violent jerk if the test needed to be stopped in the middle of a run.

Over the full 800-second cycle, the output power averaged 2 Watts and the voltage about $10V_{RMS}$. Those results are computed by considering the data as a whole. If we look at the 8-second section individually, the highest average is 14 Watts for the power, and around 26.5 V_{RMS} for the voltage (Figure 4.45).

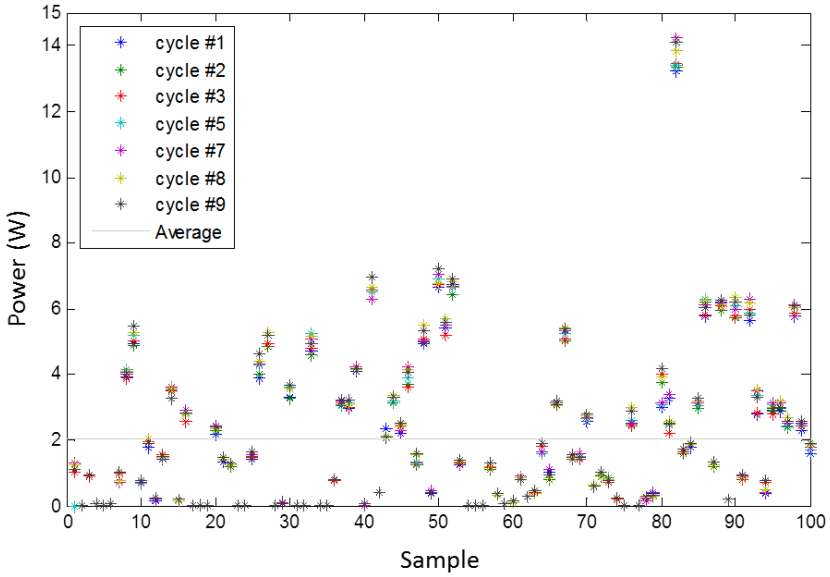

Figure 4.45: Average power for each 8-second sample

Although the average power can seem relatively low (compared to the characterization test), the instantaneous power can reach far greater values, easily in excess of 100 Watts (Figure 4.46). However, those are usually peaks of power that are unfortunately not maintained for long periods of time.

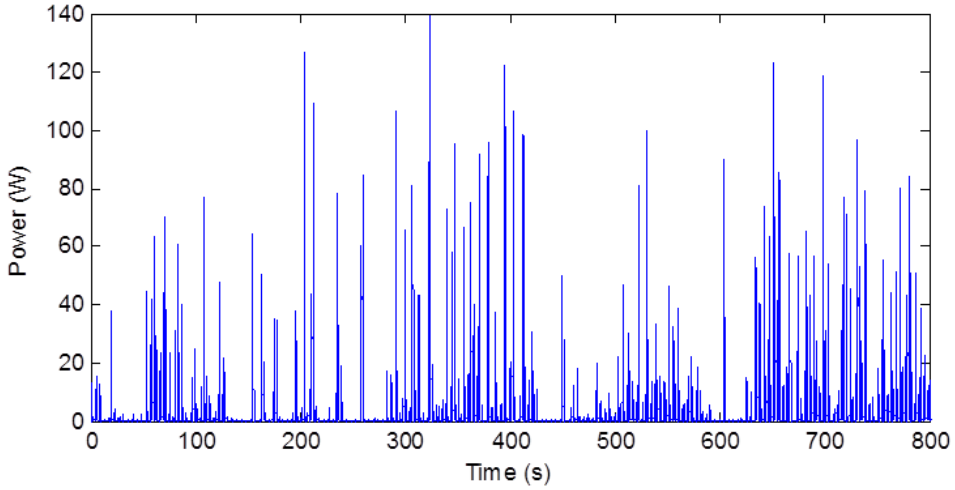

Figure 4.46: Instantaneous output power

This is representative of the nature of the field-measured displacement. Contrary to the sinusoidal input that has a constant cycle of motion, the actual suspension has significantly long periods of no activity (Figure 4.47), and mixed periods of large displacements with high accelerations. So, in some sections,

close to no output at all is generated, with the voltage barely reaching the diode threshold. Other sections manage to yield a reasonable amount of energy.

To be efficient, the system must perform well in both situations: it must be sensitive enough to harvest energy with very small inputs, and must be robust or compliant enough to withstand large, violent displacements.

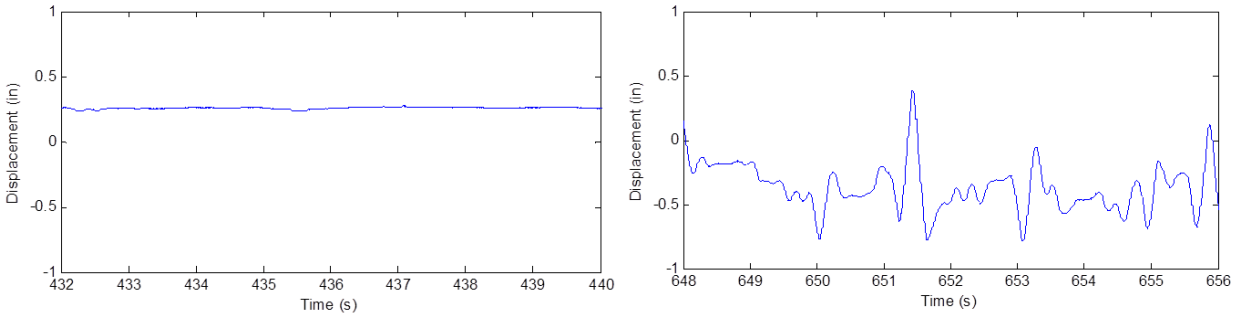

Figure 4.47: Sample with almost no motion (left), and sample leading to the best result (right)

All the results from the field data presented here were obtained with a 50Ω resistor connected as a load. Drawing more current by using a lower resistance would, up to a point, increase the output power. The system was then tested with a 20Ω resistive load. The average output power increased to 3.2 Watts, as an average over the entire suspension input. This is a decent improvement from 2 Watts, but not as much as could be expected looking at the results of the characterization tests. This can be attributed to the low pass filter that cut high velocities when the force became too high, in order to electronically protect the system. Ultimately, a mechanical torque limiter will need to be included in the prototype.

Capacitors can be placed in parallel with the load to increase the harvested energy. The idea is that the capacitor is charged during the periods when there are large motions and releases its energy when the generator output is low, filling the gaps between peaks. A test has been done with a fairly small capacitor (2350μF), yielding a nearly 10% increase in the output power, reaching an average 2.2W, without requiring noticeably higher peak forces. Larger capacitors could yield even greater improvements.

4.4.6.2 *Data from TTCi (Transportation Technology Center, Inc.)*

The Transportation Technology Center, Inc. (TTCi) is a subsidiary of the Association of American Railroads (AAR), which manages the Federal Railroad Administration's (FRA) Transportation Technology Center, located in Pueblo, Co. This large complex is dedicated to testing the reliability, wear, and fatigue of new railroad equipment such as different fasteners, new bridge design, ties, and switches. The High Tonnage Loop (HTL) is dedicated to heavy axle load; as a 2.7-mile-long loop, it consists of tangents, curves (left and right), bridges, and turnouts (Figure 4.48). TTCi has instrumented cars that run on the track and monitor the degradation of the track under heavy load, as well as the propagation of defects over time.

Figure 4.48: Satellite picture of the High Tonnage Loop

TTCi provided us with suspension displacement data measured on their instrumented car. The important difference with the data from Amsted Rail is that the conditions of the tests are known. It is now possible to associate a speed and a location with the suspension displacements, as illustrated in Figure 4.49, where the displacement is on the z-axis linked to longitudes and latitudes on the x- and y-axis, respectively. An important aspect of the data is that the displacement is contained within 0.5 inch. Therefore, the lowest and the highest positions recorded over the entire lap are less than one-half inch apart. While this difference is rather small, it is also significantly less than the range seen in the data from Amsted Rail.

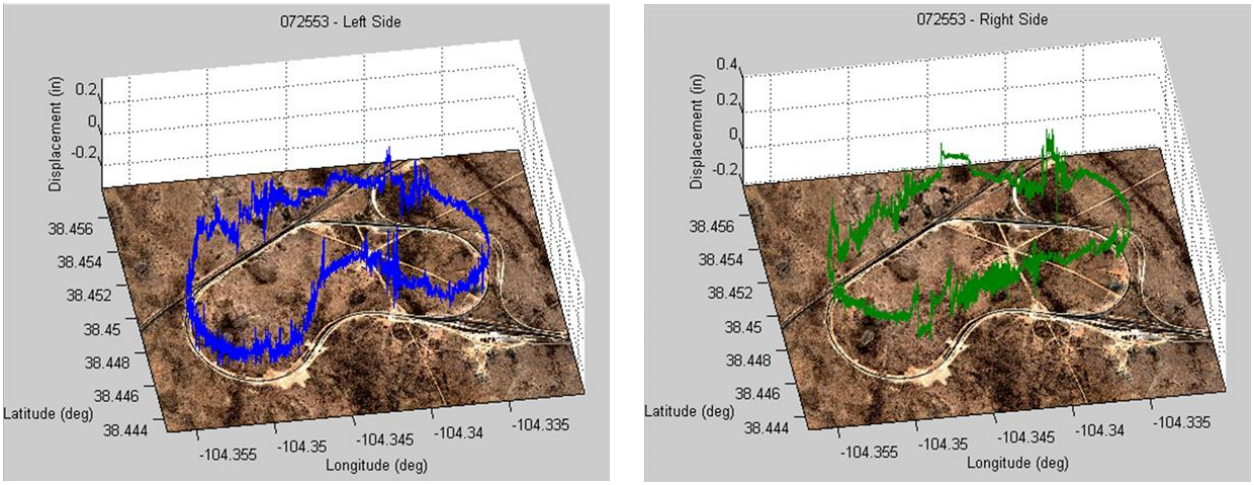

Figure 4.49: Suspension data superimposed over TTCi's High Tonnage Loop

The tests using TTCi suspension data are conducted using the same setup as for the data from Amsted Rail, with the MTS test rig coupled to a dSpace data acquisition system. A function generator is again used to provide the suspension profile. The data is also denoised by using a low pass filter. Different cut-off frequencies are used and tested.

With a 50Ω load, the system is able to produce 1.7 Watts of power, as an average over the entire loop, with the displacement measurement on the right side suspension (filtered at 15Hz). Again, although the average is not very high, the power regularly peaks at over 10 Watts, and there is a period of about 20 seconds during which an average of 20 Watts is being generated. This corresponds to the left side curve on the inside of the bean-shaped loop.

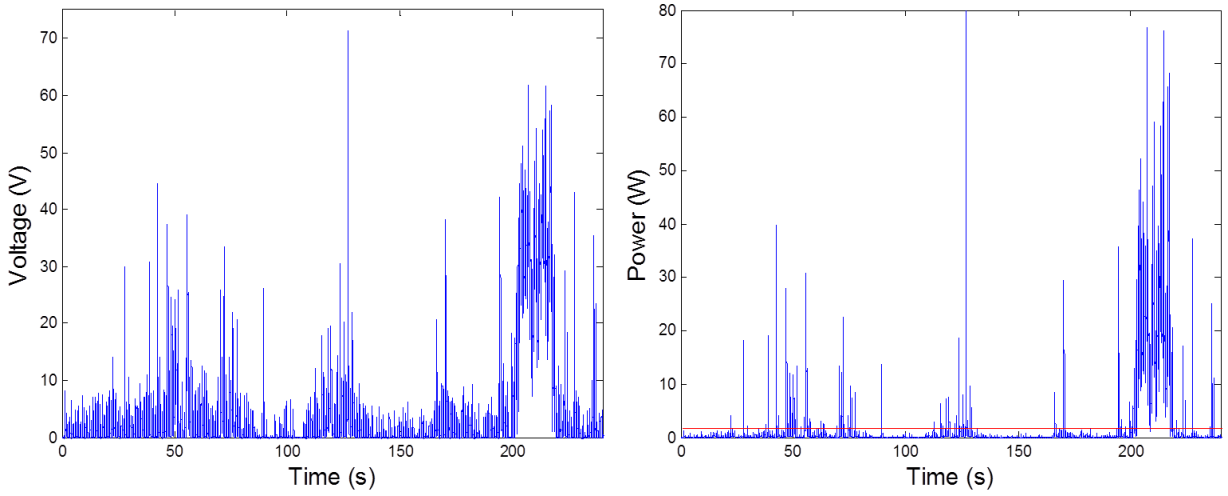

Figure 4.50: Outputs over a (replicated) lap of the High Tonnage Loop at TTC (right side)

The left-side suspension displacement only leads to about one-half watt (Figure 4.51). The power generated is more consistent than on the right side, and without the 20 seconds of "high" activity, the average remains somewhat low.

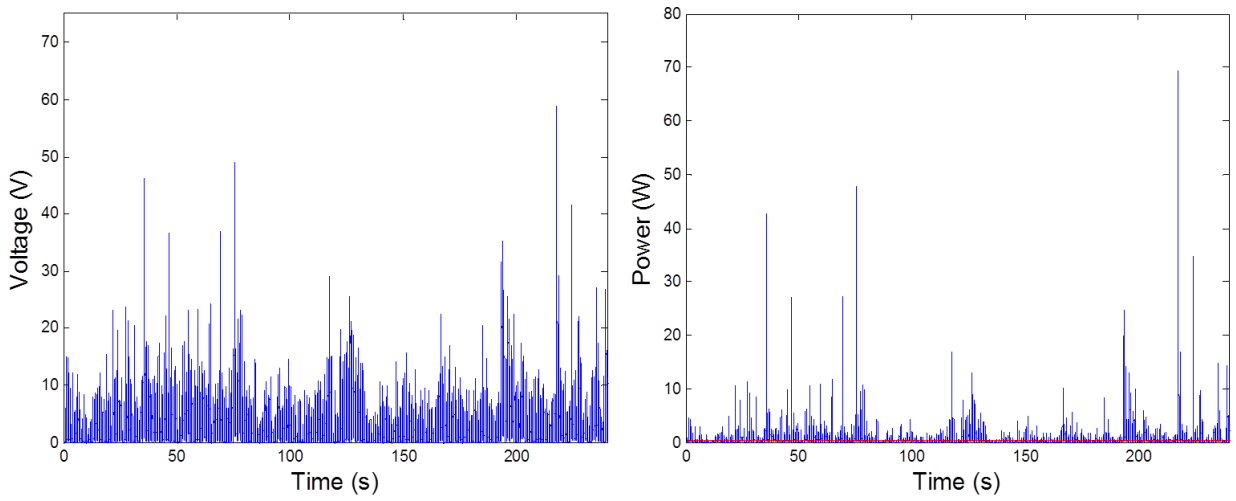

Figure 4.51: Outputs over a (replicated) lap of the High Tonnage Loop at TTC (left side)

More power cannot really be harvested as lower resistance leads, in few instances, to force too high to be safe on the ball screw. The MTS controller is set to stop the test when the forces exceed the limits. Lower resistances do constrain the controller to stop the test. This highlights that, again, a torque limiting device is needed. This would allow clearing those rare instances when the forces peak above a safe value and thus lead to more power harvested.

4.4.6.3 *Conclusion*

With actual field-measured suspension displacements, the confidence in the ability of the prototype to withstand the harsh and violent railroad environment increases. Beyond survivability, tests have proven that the current prototype is able to generate fair amounts of power from the challenging displacement sample. The prototype has been tested under those conditions for over ten hours without any issue to be reported.

It is worth mentioning that those results were obtained without modifying the system for the actual suspension data. The system remained the same it was for the sinusoidal tests. With real world data, it would be possible to improve the system to increase its harvesting potential.

4.4.7 Torque Limiter

The major needed improvement that was made clear after testing with the actual suspension data is a force- or torque-limiting device. Forces have never really been an issue with the sinusoidal inputs since they never reached, under normal conditions, the maximum rating of the ball screw. But with actual inputs, the displacements are more violent, and larger accelerations then translate to large forces due to the inertia of the numerous steel parts in rotation. One of the main difficulties in the implementation of a force-limiting device is the restricted space available, as the prototype is already densely packed.

The first idea was to use friction discs. Similar in principle to a clutch, some of those discs are connected to the ball nut subassembly, and some are connected to the part holding the one-way bearings. Although multi-disc systems are common, there is not enough space in this application to have more than two: one input and one output disc. There is no slip between the discs, and the input speed and torque match the output until the force reaches the maximum allowed by the static friction. It is then not possible to transmit larger forces. This is a simple yet effective mechanism that is used in many applications. However, the space available in the prototype is too limited to place springs capable of generating the normal load required, thus preventing the implementation of such a system.

The chosen solution is then to use spring-loaded pins that are forced into conical recesses (Figure 4.52), similarly to a ball detent mechanism. When excessive force is applied on the ball screw, the pins are pushed out of those recesses and the two sides of the torque limiter are disconnected. Only limited force can be transmitted before the pins reengage. This is a major difference from the friction limiter where, when slipping, a good amount of force can still be transmitted due to dynamic friction.

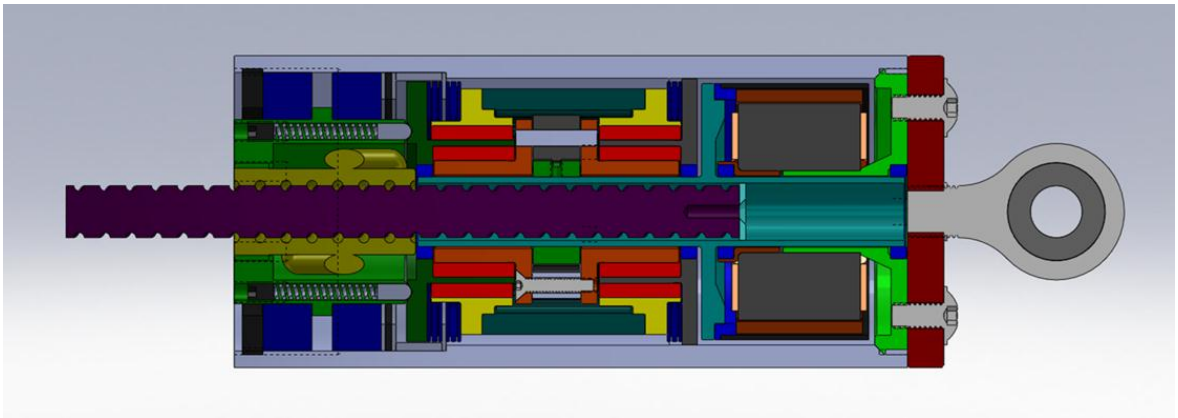

Figure 4.52: Modified prototype with torque limiter – Initial concept

The ball screw is rated from maximum forces of 800-1000lbf. With a lead of 0.5inch, this translates to 6.635 ft-lbs of torque (neglecting friction). Assuming a 120° cone angle for the pin seats and a coefficient of friction of steel on steel of 0.16, the required normal force is then 127lbs. With a total of 12 pins, one spring must produce 10.58lbf. Standard spring plungers with a maximum force rating of 11.1lbs are readily available and they will perfectly fit the application. The harvester is then modified to accommodate them (Figure 4.53).

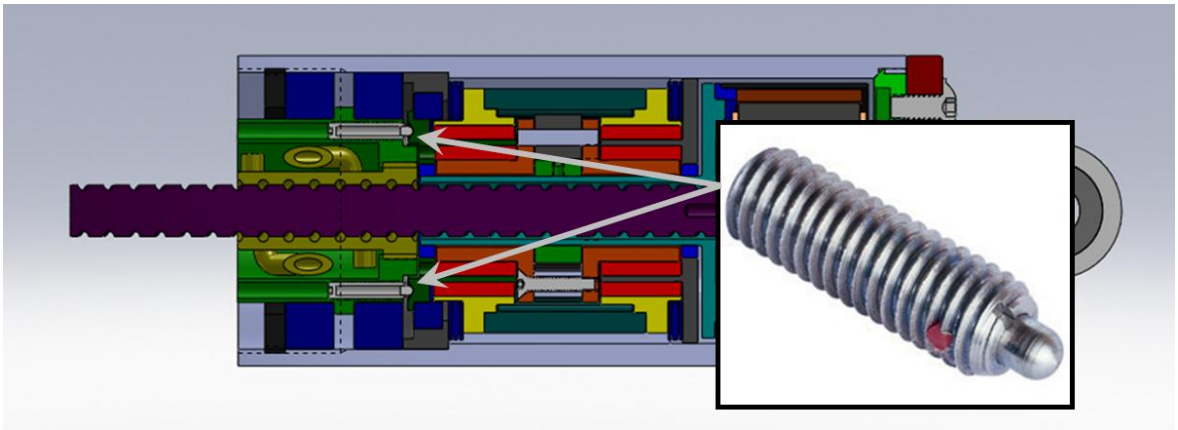

Figure 4.53: Torque limiter with spring plungers

Only three new parts are required to implement the torque limiter. A local machine shop created these and the parts were retrofitted to the existing prototype. To calibrate the torque limiter, a special adaptor plate was made (Figure 4.54 - left), so the ball nut side of the torque limiter can be spun with a wrench. The other side of the torque limiter is held in place by a shaft in which a bar has been placed to prevent rotation (Figure 4.54 - right). By pushing on a scale one foot away from the shaft, it is possible to know precisely the torque setting of the limiter.

Figure 4.54: Torque limiter calibration setup

When the torque limit is properly adjusted on the bench, the harvester is completely put back together. To properly test the limiter, the torque needs to exceed the maximum allowed. The rotor of the generator is then locked, forcing the torque limiter to trigger. On the MTS machine, a triangle displacement profile is used to repeatedly trigger the system. The maximum force recorded is relatively constant throughout (Figure 4.55), with slightly larger values in compression attributed to the increase of preload of the spring accompanying the slight deflection of the ball nut under compression.

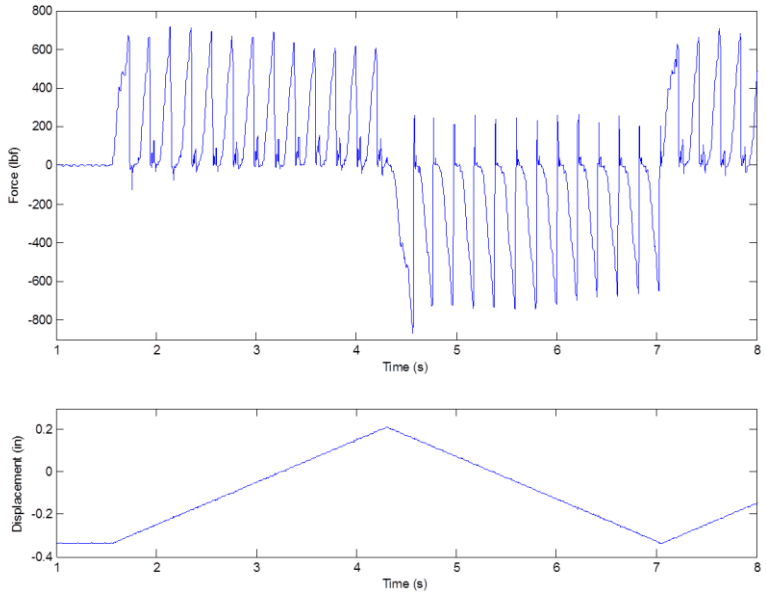

Figure 4.55: Torque limiter testing

The tests then show that this is a valid solution. However, the material used for the parts, 12L14 steel, is too soft to withstand the forces, and the pins significantly deteriorate the sliding surface (Figure 4.56). A harder steel grade, maybe quenched, is then required.

Figure 4.56: Torque limiter part showing significant wear

Several other problems are also identified. The first issue is linked to the loading. With the shape of the recesses, the pins are subjected to a lateral load, creating friction forces between the pin and the plunger body. This increases the force required to push the pin in, and thus the allowable torque. Also, it has been noted that numerous spring plungers are jamming when tested by hand before even being mounted. The solution is to replace the spring plungers by an arrangement of springs and bearing balls. The bearing balls sit in the recesses, maintained in place by springs that are preloaded with a headless screw. This has several advantages. First, the ball can be larger (twice the diameter) than the pin that has to fit in its plunger body, thus leading to lower contact pressure. Also, the ball can roll and cannot jam and lock the limiter. This also implies that the normal load will need to be slightly higher, which is not an issue since, without a plunger body, the spring can be larger. The last advantage is that this solution is significantly more economical than using spring plungers.

Another issue is that the recesses are significantly too large. Although they seem adequate in the CAD model, they are clearly oversize in comparison to the size of the pins. This was aggravated by the change from large round pins to commercially available spring plungers whose noses must have a smaller diameter to fit in the body. Different diameters are investigated (Figure 4.57), and using smaller diameters allows doubling of the number of recesses, which will lead to an earlier reengagement of the limiter.

Figure 4.57: Different torque limiter configurations

Different grades of steel are considered, and the decision was to use 8620 steel. Naturally hard and wear-resistance, it has the right properties for this application. Extensive testing on the MTS proved that this choice is adequate. The part shows minimal signs of wear after 20,000 triggering events (Figure 4.58).

The relatively large size of the bearing balls used for the limiter can also be observed in Figure 4.58 (left picture).

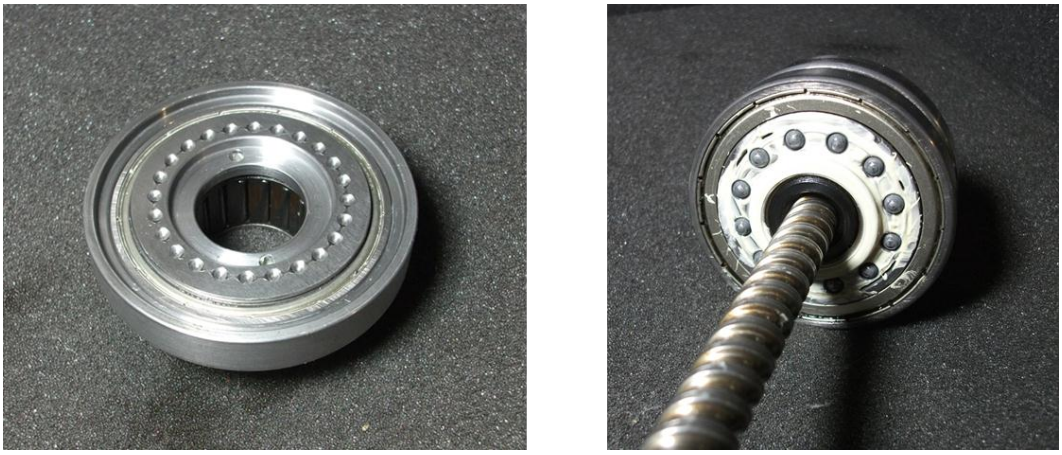

Figure 4.58: Torque limiter after testing

The right combination of elements has been found for the torque limiter which consistently limits the force below the maximum allowable by the ball screw. This feature will be included in the next generation of prototypes.

4.4.8 Conclusion

The harvester developed for a D5 spring was built on the experience gained with the first design. It proved that large amounts of power can be generated, and quite efficiently. It also demonstrated that it can withstand the harsh conditions of a freight car suspension, for extensive periods of time. Actual suspension displacements, replicated in the lab, show that the input is more random and more violent. This leads to the need for a torque-limiting system, which was consequently designed and tested.

4.6 Field Testing Prototypes

4.6.1 Introduction

The ultimate goal of the project is to develop prototypes to be tested on actual trains, and in real conditions. All the tests done in the laboratory aim at giving the highest confidence in the system. With the promising results of the harvester and the development of the torque limiter completed, the design is ready for the next phase: building prototypes for field testing.

4.6.2 Design Modifications

The new prototypes will obviously be closely based on the previous generation (Figure 4.61). However, some small modifications and upgrades are brought to the design. The first one concerns the torque limiter. Its benefits have been detailed above, and it has been shown how it is incorporated in the design. But since it's an afterthought (a late addition), the torque limiter is not fully integrated into the harvester. So, small modifications are made to make it fit perfectly in the prototype: a shoulder in the outer tube is moved, the closing nut is reduced in thickness, and threads are shortened. These are small but important modifications that bring the entire system together.

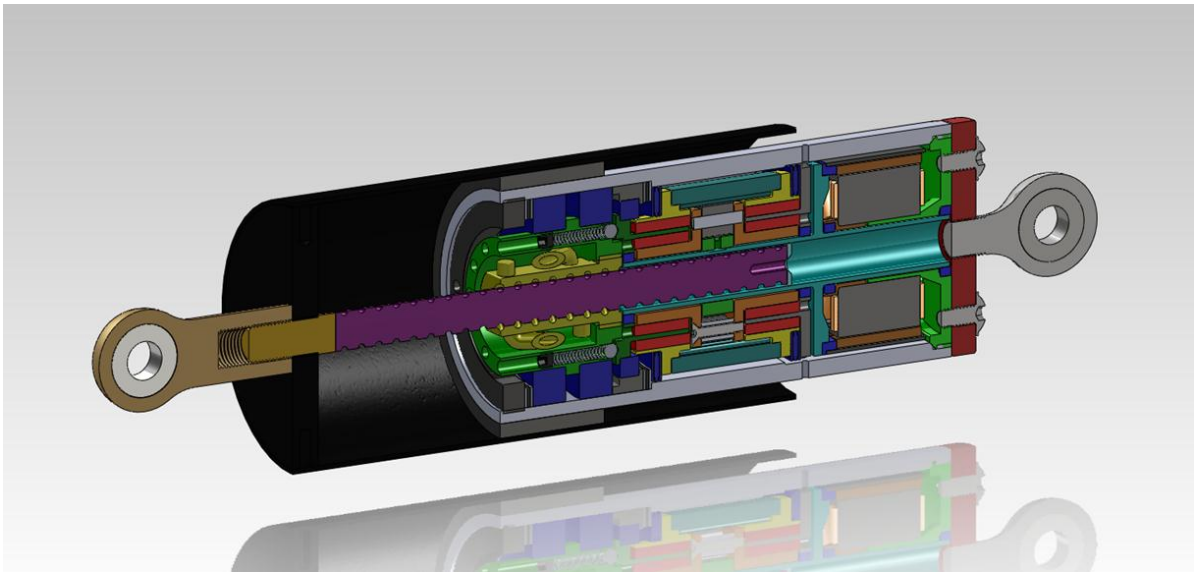

Figure 4.59: Final harvester design

Some ripples can be observed in the time plots of the force and the output voltage. This is attributed to imperfections in the machining and assembly of the harvesters. This is particularly present at the bearings holding the ball nut adapter part. Tests have shown that adding a compliant element, an O-ring, between a bearing and its support help reduce these ripples (Figure 4.60). In the tests for which results are presented in Figure 4.60, only one O-ring is used. This is why the improvement on the force ripples is mainly seen in compression. It also slightly helps in traction, by smoothing the second half of that cycle. The effect is also visible in the output voltage. O-rings have been used to improve the last prototype. They are not ideal; being round, they tend to move out of place. Therefore, they will be replaced in the new harvesters by rubber "washers." A thin (0.0625in) ring of rubber is glued between two steel washers. The latter gives the strength and stability needed while the former gives the compliance.

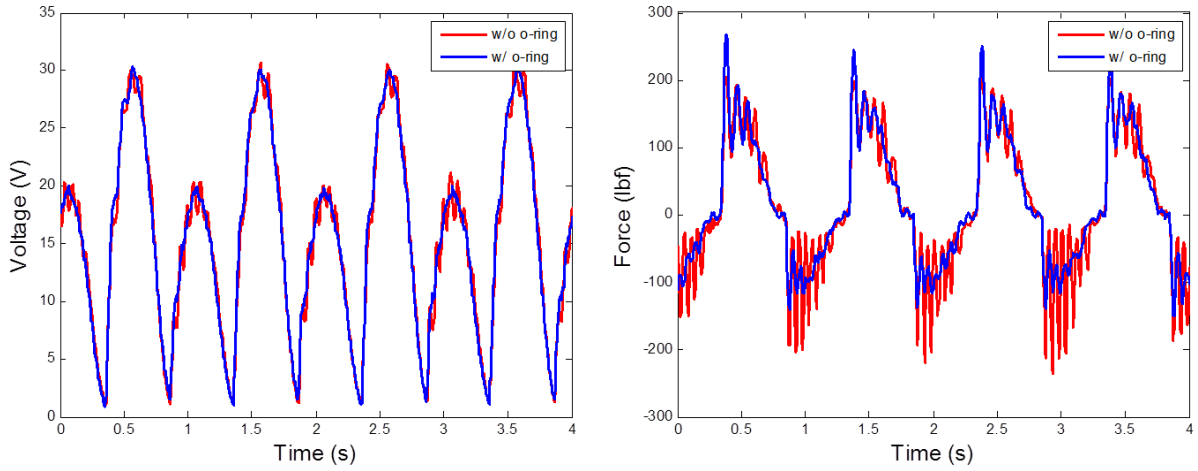

Figure 4.60: Impact of adding a compliant element

Finally, the large bearings supporting the ball nut adapter part were put in place with slight press-fit. This is suspected to slightly warp the inner race of the bearings, contributing to those ripples. In the new prototypes, they will be mounted with a tight sliding fit.

A concern is the exposure of the ball screw to dust and other contaminants. In the lab, this is a minor issue but in the field, it becomes an absolute necessity to protect the ball screw. Bellows have been considered, but could not fit in the space available. A concept using a dust cover is then adopted. It consists of a tube closed at the top and attached at the top of the ball screw. Since air will have to move in and out as that cover moves, an air filter is placed at the top of the harvester, thereby preventing contaminants from coming to the top of the system where the ball screw is located.

4.6.4 Prototypes

With the design fully defined, the next logical step is to build the prototypes. The decision was made to build eight of them. The goal was to be able to test on different platforms. The main option was to go to the Transportation Technology Center (TTC). They have been a partner on the project since the beginning, and the data that would be collected can be compared to the results of the lab tests and the conditions on the High Tonnage Loop (HTL) which are known and controlled. Another option was to work with railroad companies like Norfolk Southern or Union Pacific to mount the harvester on one of their cars and allow it to run in revenue service. Those companies have already expressed their interest for testing the technology. So, in order to take advantage of those testing opportunities, it was important to have several prototypes, and eight seemed to be a good quantity for a field testing campaign.

The manufacturing of the custom parts, just as for the previous generation of prototypes, was done by a machine shop. The laboratory is not equipped with the machining equipment required for the level of complexity of the parts, even if they are rather simple. The sheer number of parts (Figure 4.61) also makes CNC machining an obligation.

Figure 4.61: Custom parts for eight harvesters

In Figure 4.61, it is worth mentioning the presence of two D5 springs in the background, which gives an idea of the size of the suspension springs of a freight car. Usually eight or nine of them are used at each corner of the car.

The rest of the parts which includes generators, ball screw/nut, rod ends, springs, hardware was ordered from various suppliers. Some of the parts were then customized in-house (Figure 4.62). For example, the ball screws were cut to size and tapped at both ends, the rotors of the generator were cut to fit the custom shaft, and the diameter of the sun gear was increased. The dust covers were fully machined in the laboratory. Also, just as for the previous generations, the windings of the generators were completely redone to fit the application.

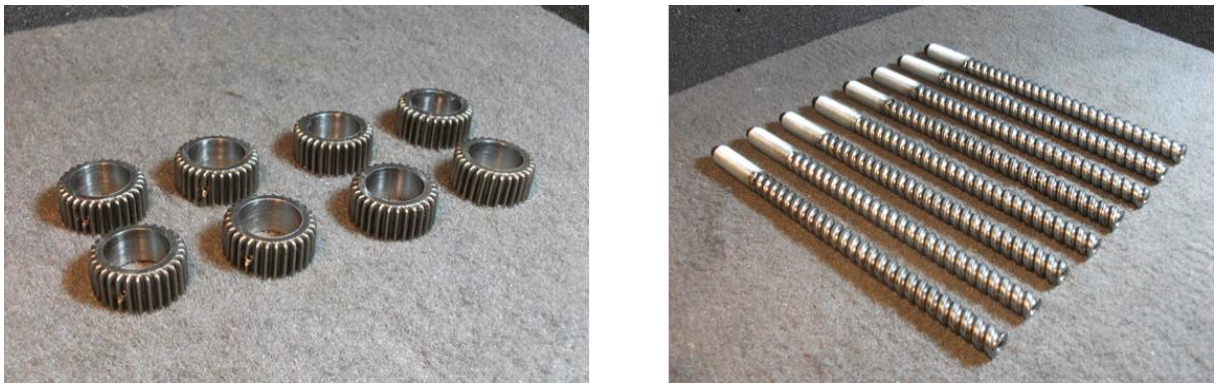

Figure 4.62: Part of the modified standard parts

Some of the custom parts had to be reworked as they were at the extreme of the tolerances, and because some of the bearings turn out to be slightly lower than in the previous prototype, even though they were ordered from the same suppliers. Small tweaking was required for most of the parts in order to assure proper fit.

Once all the parts were prepped, the harvesters were assembled and painted (Figure 4.63). Three-prong connectors were fitted to the phase leads, which allow an easy connection with the electronic box. With the cover removed, the air filter is clearly visible at the top of the harvester tube.

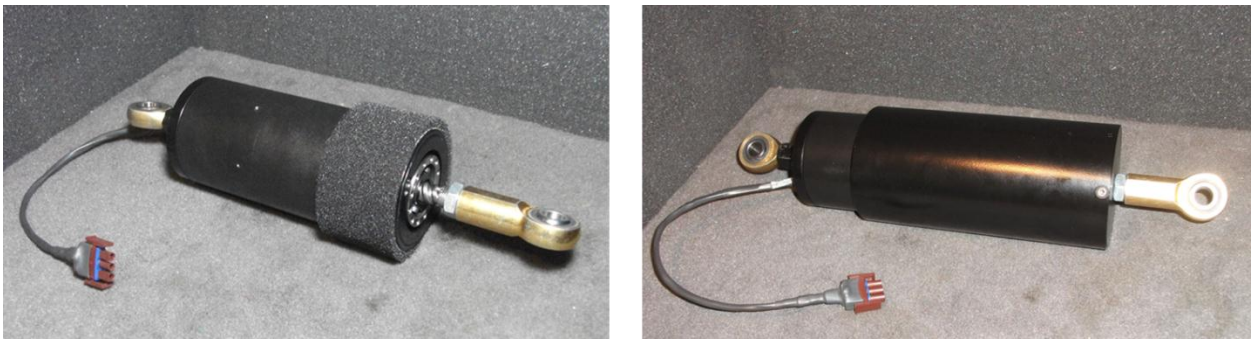

Figure 4.63: Fully assembled harvesters with (left) and without (right) the dust cover

At this stage, the harvesters are almost ready for field testing. However, they must be tested one last time on the MTS test rig for an ultimate inspection.

4.6.5 Lab Testing

Once assembled, all the prototypes were tested on the MTS test rig, both for characterization with short sinusoidal tests and for durability (100,000 cycles). The results of the sinusoidal tests, as can be expected, were close to those of the previous generation. Again, the average voltages are, for a given resistor, proportional to the speed or to the amplitude, since the tests are all done at 1Hz (Figure 4.64). This leads to the average power being quadratic.

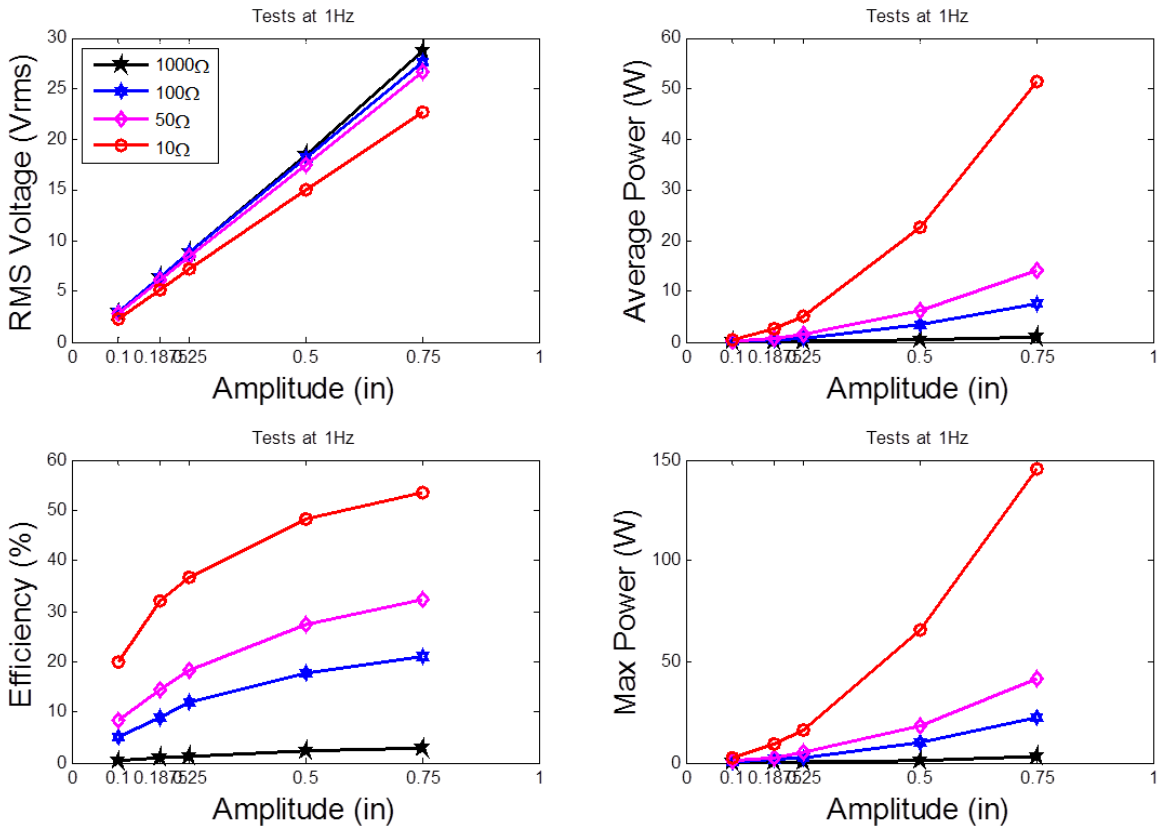

Figure 4.64: Characterization test results of the harvesters for field testing

The resistive load connected to the harvester is directly linked to the amount of power being generated. Another important electrical parameter is the capacitance. A capacitor can be connected after the diode bridge to smooth out the output voltage, filling the gaps between voltage peaks. Numerous combinations of resistors and capacitors were tested (Figure 4.65).

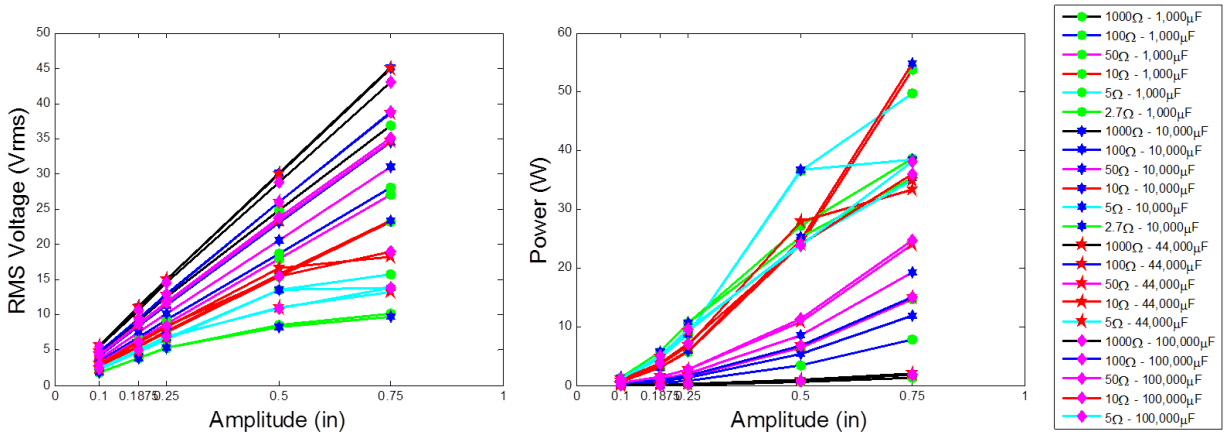

Figure 4.65: Tests of various resistor/capacitor combinations

One would assume that the larger the capacitor, the better: a lot of energy can be stored, and the power would be spread out over a long time. In some cases, this can quite well describe the situation. For example, with a 100Ω resistor (Figure 4.66 - left), larger capacitors lead to larger average voltages and powers, up to a certain point where the harvester is no longer capable of generating larger outputs. In Figure 4.66, with a 100Ω resistor, it can be observed that a 44,000μF capacitor leads to better performance than a 10,000μF capacitor, which does better than a 1,000μF capacitor. But a 100,000μF capacitor does not bring any improvement over the 44,000μF capacitor: the harvester has already given its maximum.

With larger capacitors, an issue is that they will need to draw a lot of current to become charged. Currents translate to torque and then to forces. If they are too high, they lead to torque exceeding the maximum allowable by the limiter, which then starts slipping. This can be seen in the plots on the left side of Figure 4.66. At 0.75-inch amplitude, the larger capacitors lead to lower power harvested because the current drawn triggers the torque limiter. In this case, a smaller capacitor is then desired.

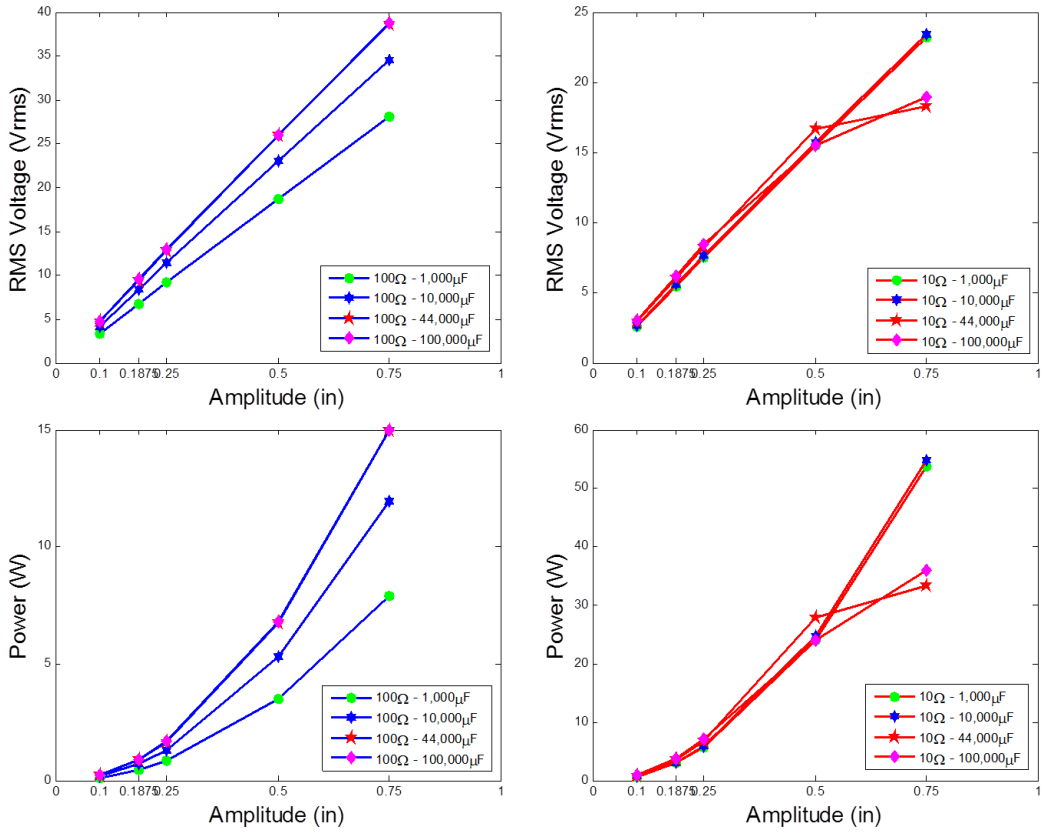

Figure 4.66: Average voltages and powers for various capacitors

An important parameter of the generator is the Kv value, which is the ratio between speed in rpm and voltage. Before soldering the connector to the generator wires, it is possible to test different configurations. One configuration consists of dividing the strands into two cables and connecting them in series. This doubles the effective number of turns per arm of the generator. This should also double the electromotive force of the generator, and potentially quadruple the output power. One drawback is that since the space is limited, the number of wires cannot increase, and the internal resistance of the generator quadruples. This in turn increases the copper losses, but the idea is that the performance would increase as the generator can operate in a better, more efficient range. Tests clearly show that the open-circuit voltages (or with high resistors) double (Figure 4.67). As the resistance decreases, the losses increase and the values are no longer quite doubled. But the power nevertheless greatly increases for a given resistor. It is clear in Figure 4.67 that the results with a 10Ω resistor are not quite as good. The first three points on the voltage plots are aligned and follow the expected linear relationship. However, starting with the fourth data point, voltages and powers are not where they would be expected. This is due to too much current

drawn at peak voltage, leading to the torque limiter slipping as explained for the capacitors. So, a lower Kv value can lead to more power being harvested, given that it does not cause the limiter to trigger.

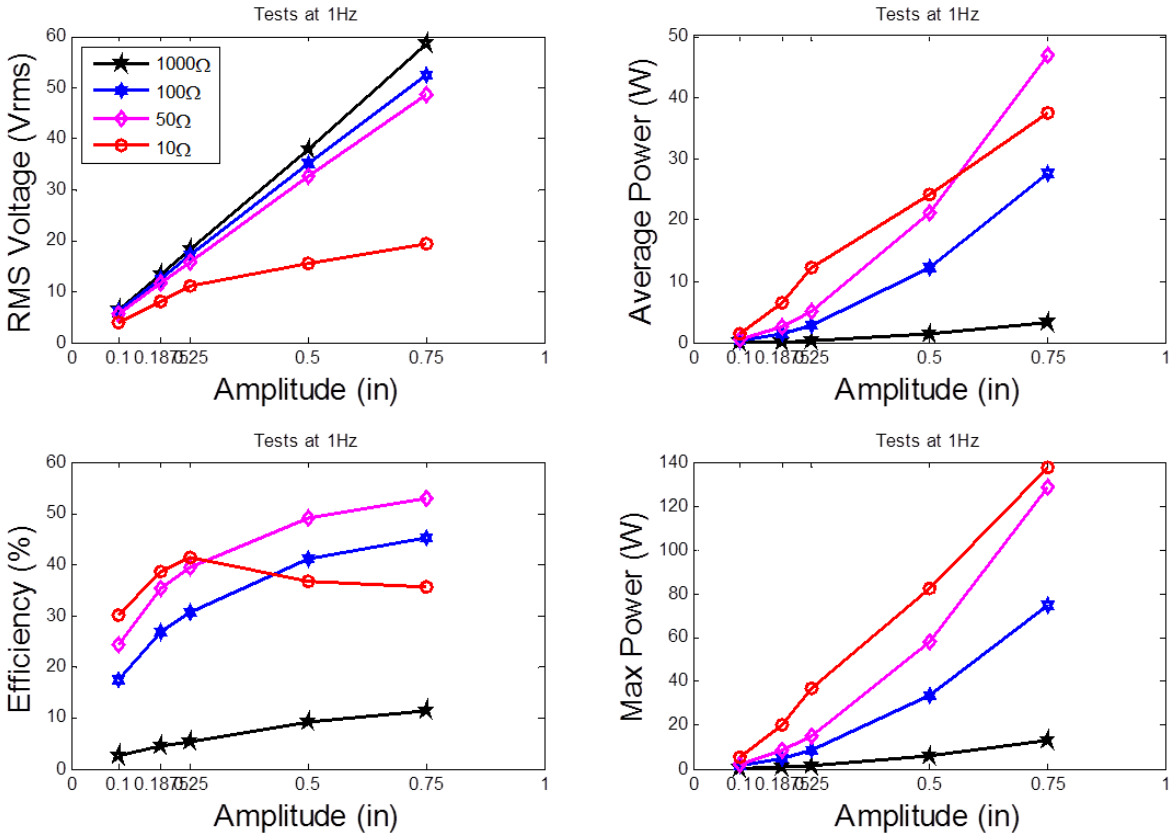

Figure 4.67: Characterization test results with a lower Kv

The previous discussion regarding the resistors and capacitors is based on sinusoidal input tests. It now must be determined how this correlates to more realistic displacements. So, the data provided by TTCi is used to test different combinations. For larger resistors, a larger capacitor (22,000μF) provides better performance (Figure 4.68). But at 10Ω and below, a 1,000μF capacitor will lead to more power being generated. As explained before, this is due to the torque limiter being triggered. So, just as for sinusoidal motions, the capacitor has to be chosen taking into account the resistive load.

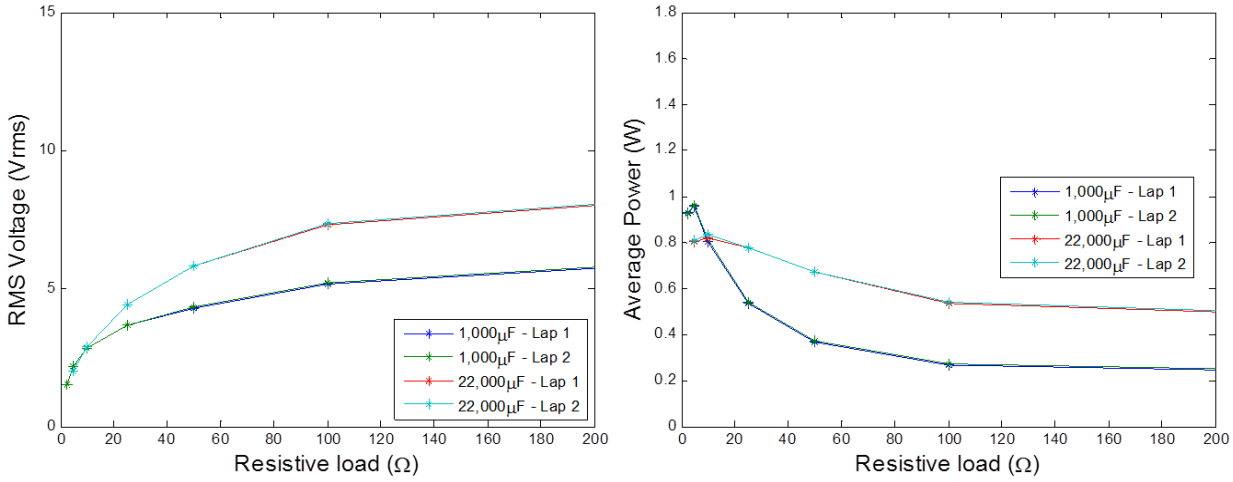

Figure 4.68: Test results for different resistor/capacitor combinations using TTCi's data (left side)

In Figure 4.68, the average power is fairly low compared to the numbers seen with sinusoidal inputs. This is due to the fact that the displacement measured is fairly small and highly inconsistent. The time plot of the measured output (Figure 4.69) clearly shows that significant voltages are produced, and that power is consistently above 10 Watts. But those values are unfortunately not maintained for very long, and the overall average is severely affected by these periods of very low power generation.

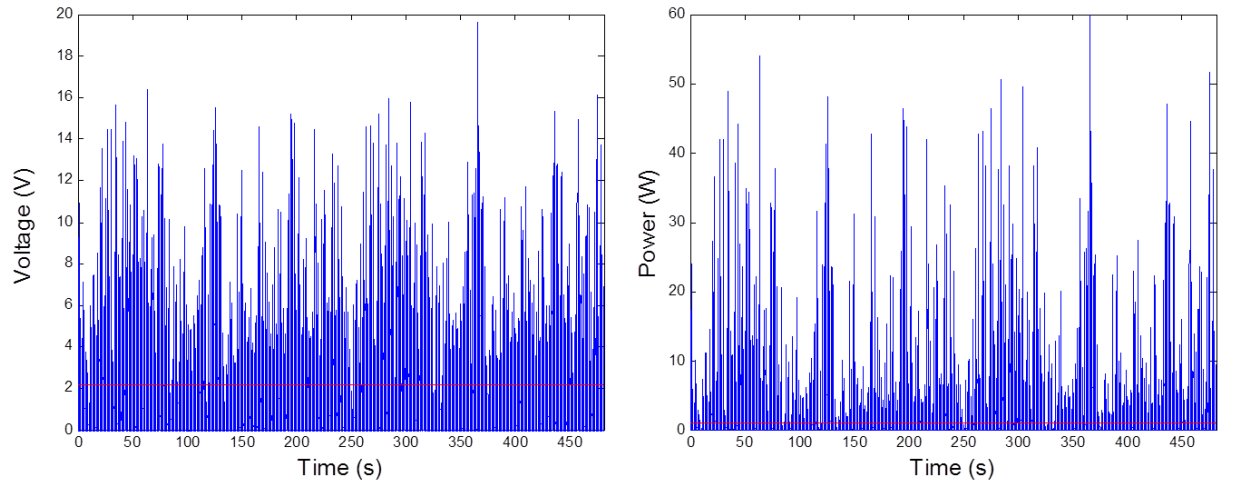

Figure 4.69: Time output with a 5Ω resistor and a 1,000μF capacitor (left side - 2 laps)

It has been mentioned earlier that the Kv value of the generator has a direct impact on the voltages generated and thus the power. The same idea of placing half the strands in series has been tested with the suspension displacement measurements provided by TTCi. The improvements are quite noticeable

(Figure 4.70), with the power more than tripling for some values of resistors. With low resistance (below 10Ω), the benefits are not as great, and the limiter gets triggered frequently. Depending on the load, this can be a great way to improve the power generated.

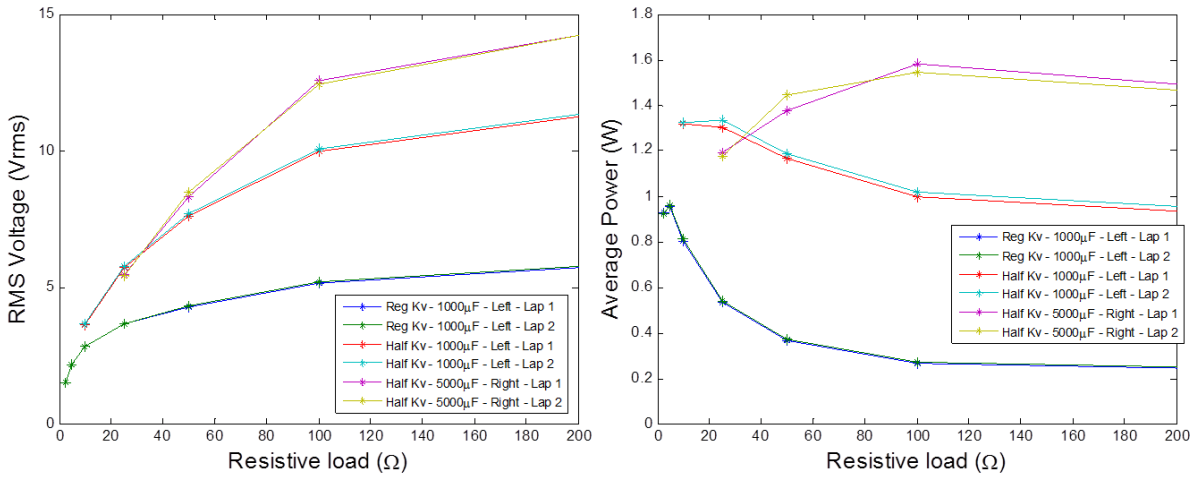

Figure 4.70: Test results for different Kv values

As mentioned before, the measured displacement is very limited. Over a complete lap of 2.7 miles lasting about 4 minutes, the suspension remains within one-half inch. Tests are then conducted using displacements equal to twice the values recorded. This can give an interesting insight into what would happen if the displacement has been underestimated or if the suspension happens to move more in a different situation, like in revenue service. Although it is two times larger, the peak-to-peak displacement remains within one inch. Doubling the displacement also means doubling the velocity across the suspension, which should lead to twice the output voltage. This would then mean four times the power. In addition, since the generator is allowed to spin on its own inertia, the output power could be even higher. Tests prove that in these conditions, it is possible to get over 4 Watts of power, which is more than four times what was generated previously. The profile of the output voltage remains similar with a lot of peaks, but it is more uniform (Figure 4.71), now consistently above 8 Volts. With the quadratic relationship, the power is a little more spread out, but it remains high quite regularly.

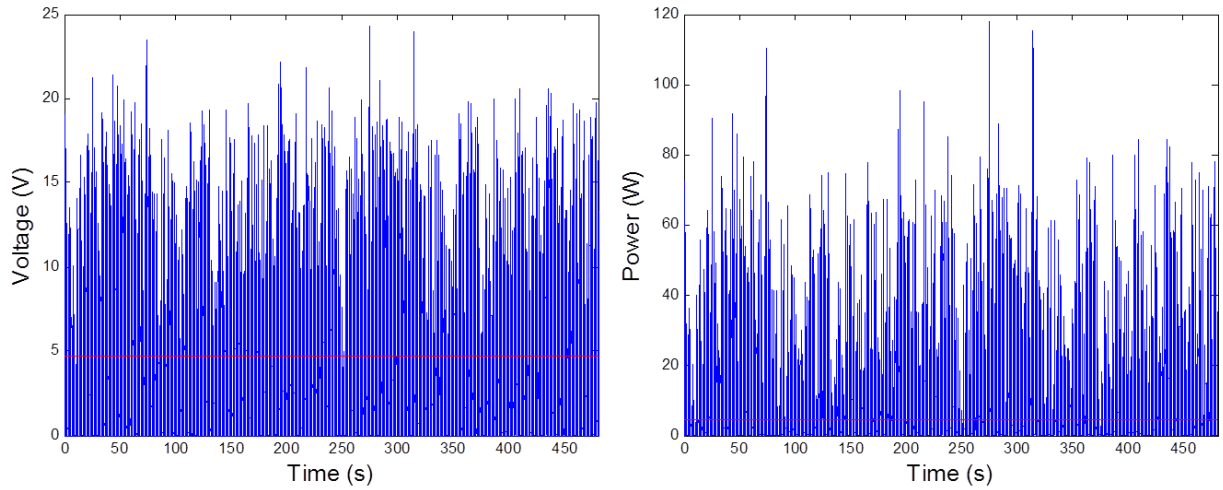

Figure 4.71: Time output for twice the suspension displacement (5Ω resistor - 1,000μF capacitor)

Numerous combinations of resistors and capacitors are also been tested using doubled displacements. Similarly to the previous results, for some values of resistor (those at the lower end), adding a larger capacitance is detrimental (Figure 4.72). But above 50Ω, it leads to larger amounts of power being generated. This is clearly illustrated on the plot of the voltage or power versus capacitance in Figure 4.72. For some resistances, the curves are monotonically increasing, proving that larger capacitors improve the performance. Conversely, some curves are decreasing with increased values of capacitance, proving that the capacitor should be rather small. Again, the capacitor has to be carefully selected, taking into consideration the value of the resistive load that will be in the circuit.

Testing with doubled displacements also proves that the harvesters are capable of withstanding harsher conditions without any incident.

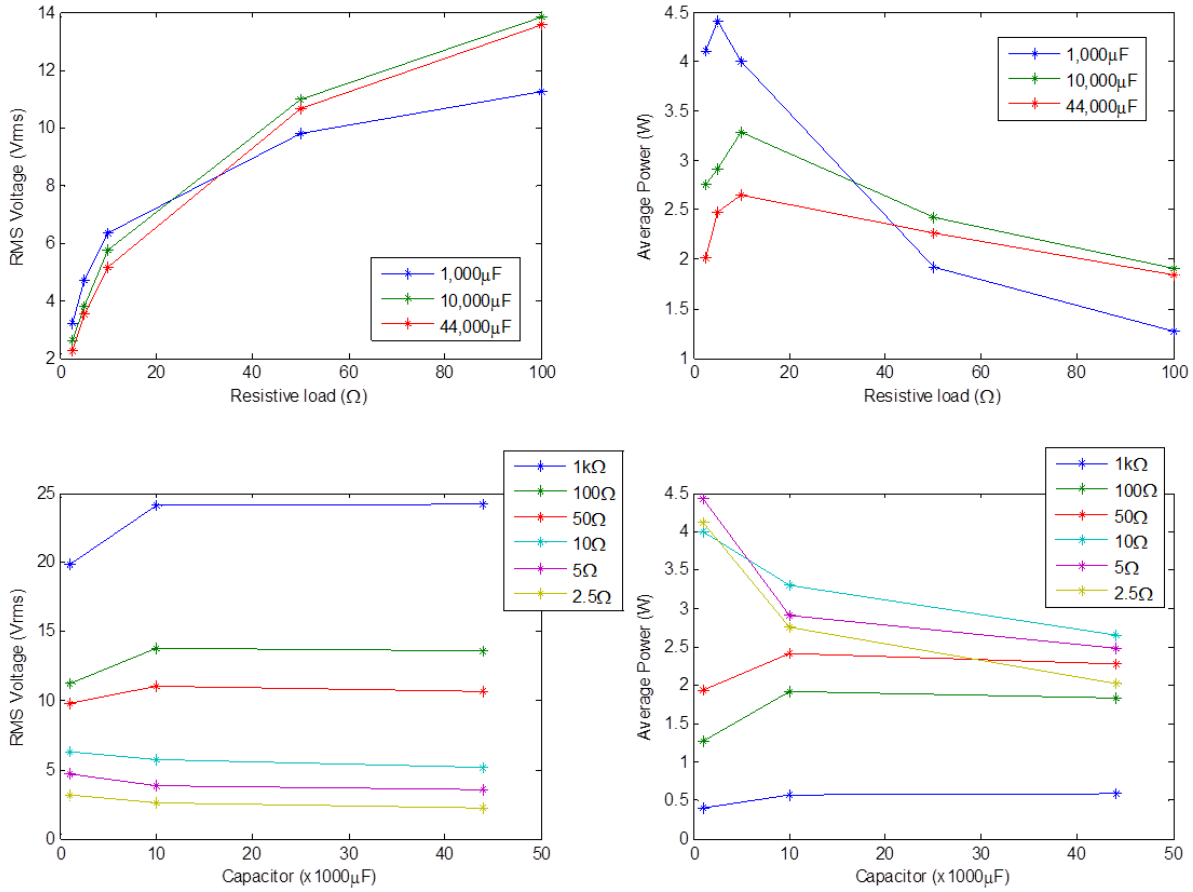

Figure 4.72: Test results for various resistors/capacitors using twice the suspension displacement

The harvesting systems prove that they can produce large amounts of power in a consistent fashion with regular inputs like the sinusoidal displacements, or in a more erratic way with the highly variable displacement seen across a freight car suspension. In the latter case, the average power harvested suffered greatly from the periods of time with low activity. It remains that few watts of power can be expected, and that tuning after collecting data in field testing could bring out the full potential of the system.

4.6.7 Deployment at the Transportations Technology Center

At the end of October 2012, TTCi permitted access to their facility and allowed two harvesters to be mounted on their instrumented freight car (IFC).

4.6.7.1 *Mounting System*

The harvesters are designed to be placed inside the suspension spring. However, for the preliminary tests, it was immensely more practical to mount them on the side of the truck. To get the harvesters inside the spring nest, it would require lifting the car body, taking the bogie apart, setting the harvester, and putting everything back together. However, this is too work intensive, especially for a system that has not yet been proven in the field. So, the solution for the test campaign was to mount the harvester along the springs, in parallel with the suspension, right on the side of the truck (Figure 4.73).

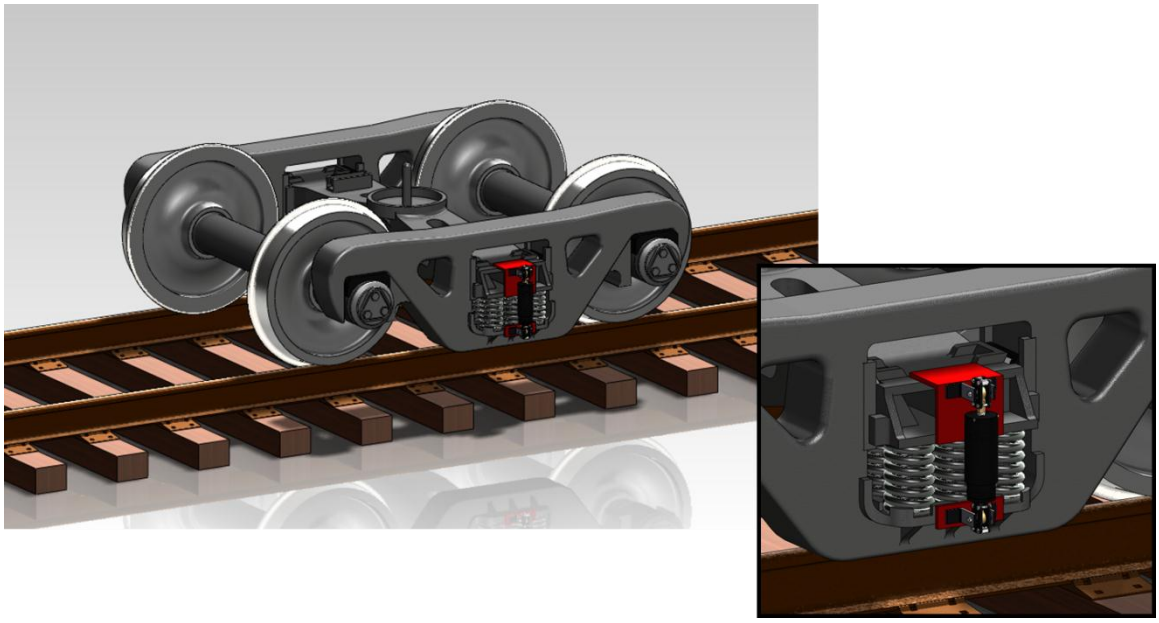

Figure 4.73: Design of the mounting system for field testing

In Figure 4.73, the red brackets are the parts that need to be welded at the end of the bolster and the side of the side frame to provide a flat surface on which the harvester will be mounted. These red parts were the responsibility of TTCi, and all the parts to be mounted on these brackets were provided, including the harvester itself, as well as the angle brackets and the anti-swivel mechanism.

As forces are applied to the harvester, the ball screw will make the ball nut rotate and, as a results a torque is created. The mounting system needs to be able to counter that torque and prevent the harvester from twisting around the longitudinal axis. Any motion allowed around that direction is that much less motion

generated at the ball nut, and thus it is a loss of potential for harvesting energy. So, it is crucial to limit the motion of the harvester around that longitudinal axis as much as possible. The ball joints allow three degrees of freedom in rotation. It is desirable to keep two of them, as they can compensate for any misalignment. The degree of freedom around the longitudinal axis of the harvester, as explained just before, must be eliminated. After considering several designs, the chosen solution is to use guide plates on both sides of the ball joint, with an oblong hole that allows the two desired rotations and blocks the third rotation (Figure 4.74). The shaft connected to the ball of the rod end can rotate freely around its axis, sliding up and down in the oblong shape to allow a second rotation, but the first degree of freedom is removed as the shaft is not free to move, hitting the edge of the oblong hole.

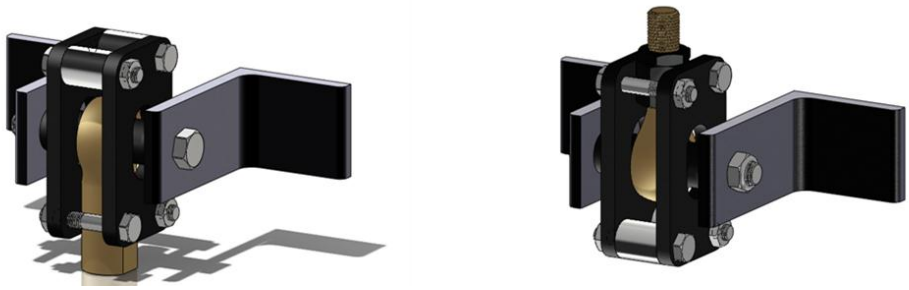

Figure 4.74: Anti-swivel mounting system

The guide plates are manufactured by a waterjet company. Two step sleeves that fit on each side of the ball joints are also required. The rest of the parts are common hardware elements. With the anti-swivel mechanisms mounted, the harvesters (Figure 4.75) are ready for testing on a freight car.

Figure 4.75: Complete harvesters, ready to be mounted on a train

4.6.7.2 *Electronics*

The goal of the first test is to determine the power generated. The energy is then not actually harvested but is sent to a resistor and dissipated as heat. The electric circuit is similar to that used during lab tests (Figure 4.36), but is packaged to be mounted on a freight car (Figure 4.76). The output of the box is connected to a power resistor. It is not incorporated in the electronics box so it can easily be changed from one test to the next.

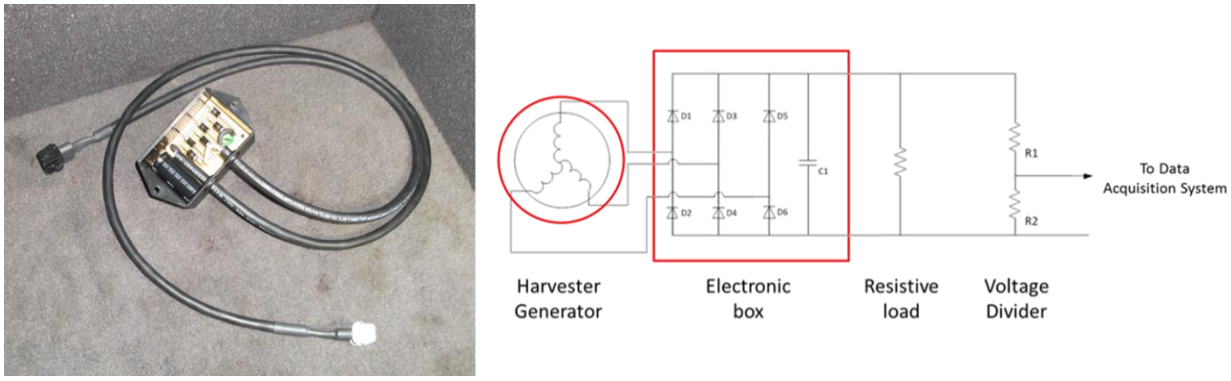

Figure 4.76: Electronic box and electrical schematics

A joint effort with a team from the Electrical Engineering department focused on providing a smart battery charger capable of taking the random output of the harvester and conditioning the power into a form suitable for the battery. The circuit acts as a 10Ω resistor, and efficiently charges the battery with the corresponding (highly variable) power. Preliminary tests (Figure 4.77) show good results, and a second, more refined prototype is currently in development.

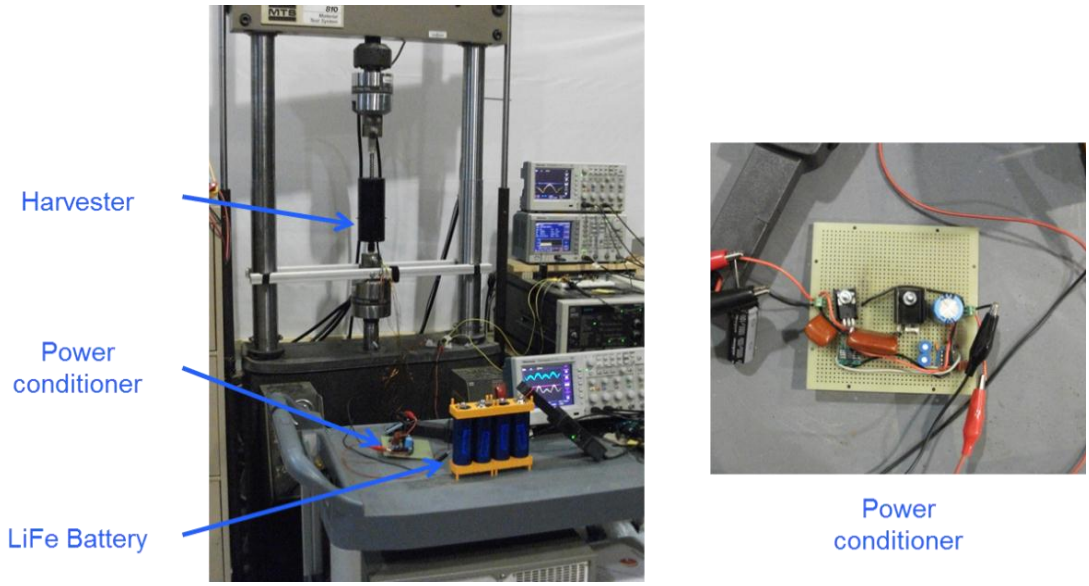

Figure 4.77: Testing of the power conditioner

After validation in lab testing, this power conditioner will join the harvester in field testing on actual freight cars, and energy will actually be harvested and stored in a battery.

4.6.7.3 *Implementation on the Instrumented Car at TTC*

As mentioned before, the Transportation Technology Center (TTC) is a large testing facility located in Pueblo, CO that is managed by a subsidiary of the American Association of Railroads (AAR). They have dedicated tracks and rolling equipment for assessing any railroad-related technology. One of their tracks is the High Tonnage Loop (HTL), which is used to test freight equipment. On the HTL, TTCi operates a train of three locomotives hauling about 100 coal cars (Figure 4.78), with the first one being instrumented. Various parameters are recorded, such as the speed of the train, its position, the accelerations on different parts of the car, or the displacements of the suspensions.

Figure 4.78: TTCi's test train

The harvesters are mounted on the instrumented car, which makes the data collection that much easier as it is possible to use two inputs of the data acquisition system already installed on the car. The brackets (those parts in red in Figure 4.73) were welded on the bogie by a machinist from TTCi (Figure 4.79), and the harvesters were in turn welded to the brackets. It is important to set the harvesters to the right height before welding, as their stroke must match that of the suspension. If set too long, the suspension would bottom out on the ball screw and certainly ruin it. If set too short, the suspension when fully extended

would force the ball screw through the ball nut, thus losing the bearing balls and rendering the system inoperative. The right length for the harvester is based on the length of the spring at the time of installation (see Appendix C).

Figure 4.79: Mounting of the harvester onto TTCi's instrumented car

Once the harvesters are welded in place, the electronic box is secured on the side frame and connected to the harvesting system (Figure 4.80). A box for the resistive load and the voltage divider is placed on the car as well, and is connected to the electronic box. The ouput of the voltage divider is recorded by the onboard data acquisition system.

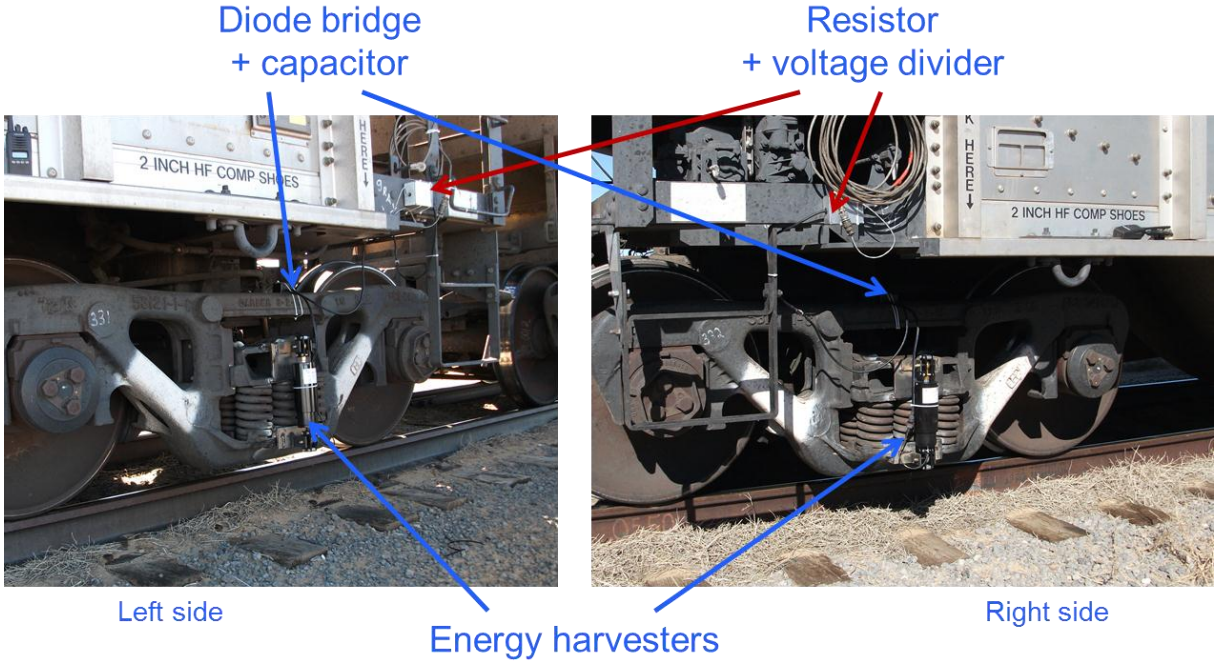

Figure 4.80: Harvesting system fully in place and ready for testing

All the elements are then in place to start testing the harvesters on a real freight car in real conditions.

4.6.7.4 *Test Results*

At TTCi, regular days of testing consist of running the train all night, thus covering about 100 laps of the HTL, although the number varies as the train sometimes has to stop for inspection or repairs. During the first three days of testing, the data acquisition system did not properly record the output of the harvesters. Therefore, there is no possibility of knowing how well the systems performed. The only information available is that the harvesters are capable of withstand the harsh conditions encountered in a freight car suspension.

Unfortunately, the problem on the DAQ system could not be resolved before the end of TTCi's testing season. Tests are run every day for a few months at a time, and are then stopped to prepare other tests and perform the maintenance of the track and train. The harvesters were installed on the last week.

Therefore, no data has yet been recorded over the entire track. But outside the testing season, small tests were also conducted. One test focused on the bridges (section circled in red in Figure 4.81). For that particular test, the train was moving back and forth on the track and was maintained at regular speed only on the bridges. This set provided much less data than what would be obtained during regular testing, but it did provide a formidable look at the performance of the harvesters.

Figure 4.81: Section of the High Tonnage Loop (HTL) under test

A 100Ω resistor was used as load. This choice was driven by the desire to not draw too much current or put too much stress on the systems for the first tests, while still getting some power out. The profile of the output voltage and power (Figure 4.82) is what can be expected: highly variable output with numerous peaks. Over the various sets of data available, the general trend is to have average voltages between 4 and

$6V_{RMS}$, leading to powers in the range of 200 - 300mW. This is in the range of what has been obtained during laboratory tests. Some recordings offer better averages, with voltages around 6 - $7V_{RMS}$, for powers of 350 - 500mW.

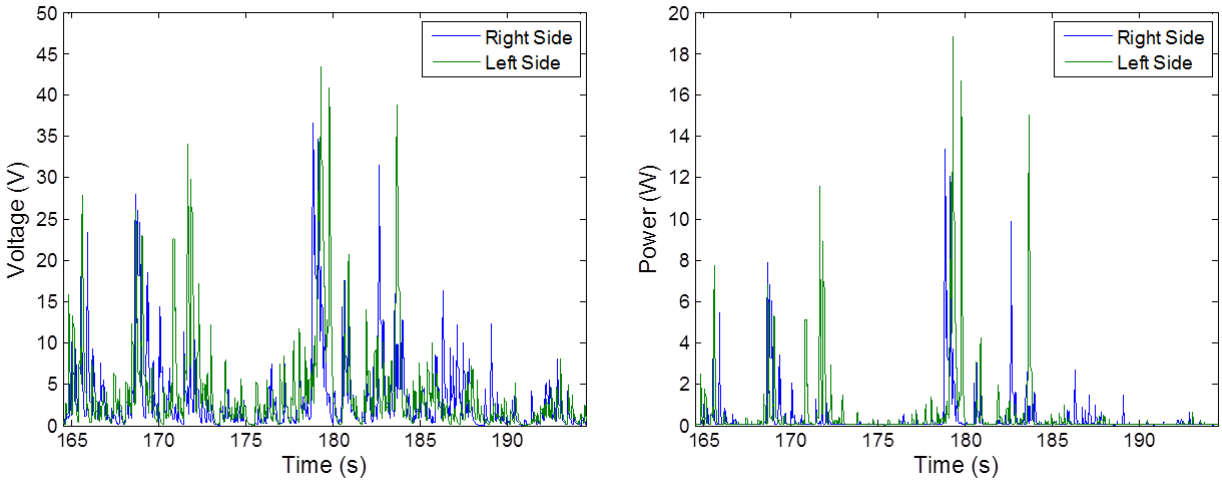

Figure 4.82: Test results from a section of the track at TTCi

Those very first results were, again, obtained with a 100Ω resistor. The power generated can be increased by drawing more current. After analysis of the field data, it will be possible to optimize the harvester parameters (for example, Kv value and capacitance), to increase the power generated.

4.6.8 Concluding Remarks

A lot of work has been dedicated to make the testing of the harvesters in the field a reality. This has finally been possible with the help of TTCi. While there is only a limited amount of data available right now, the recorded voltages and powers are still promising, even if they are not very high. Testing at TTC should resume soon and more data will be available, providing much more insight into the actual potential of the current harvesters.

4.7 Conclusion

In this chapter, the potential of prototypes using a rotating generator have been demonstrated. The ability to have a gearbox to increase the generator is the key to good performance. Several upgrades have been

made to the first prototype, showing at each step what improvement resulted. The second generation of prototypes is built on the experience gained with the first design, proving that large amounts of power can be generated, and quite efficiently. It also demonstrates that it can withstand the harsh conditions of a freight car suspension for extensive periods of time. Actual suspension displacements, replicated in the lab, show that the input is more random and more violent, leading to the need for a force-limiting system. With the torque limiter developed, a pre-production series of eight harvesters have been produced. Two of these have been deployed at the Transportation Technology Center in Colorado, and they prove that the systems are field ready and show encouraging results. More harvesters should be deployed in the near future.

This page intentionally left blank.

Chapter 5

Energy Harvester Modeling and Simulation

As a designing/optimization tool or as an element of a suspension analysis, it is important to have a numerical model of the energy harvester. Using Newton's second law and Euler's second law, the equations of motion for each sub-assembly (Figure 5.1) are determined.

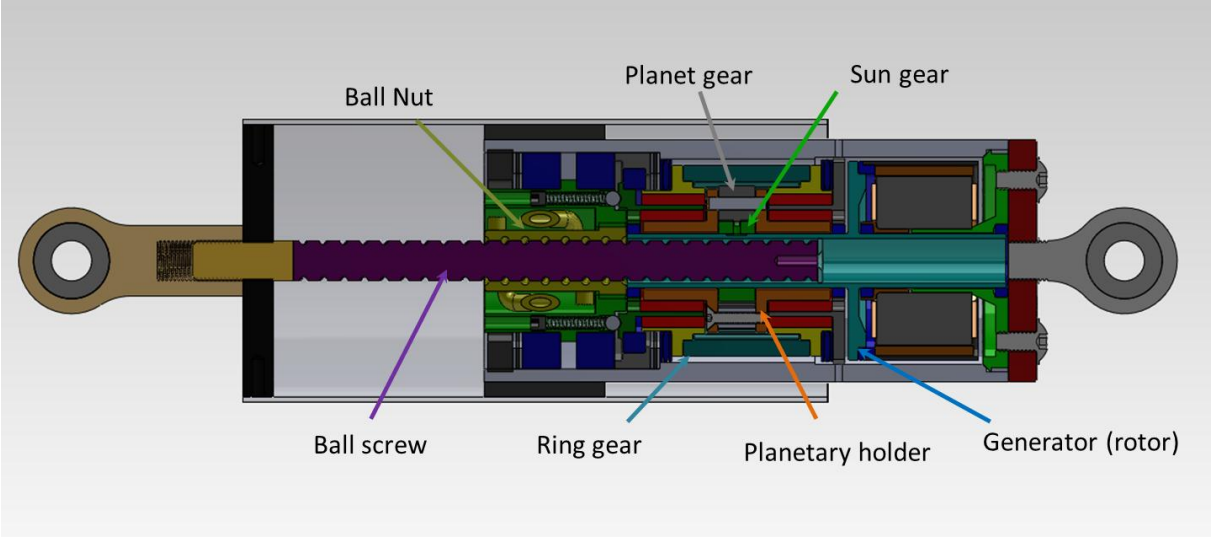

Figure 5.1: Prototype sub-assemblies

The motion of the ball screw can be described by:

$$M_S \cdot \ddot{x}_S = F_{in} - F_N \tag{5.1}$$

For the ball nut assembly, it is interesting to look at both the force in the longitudinal direction and the moment:

$$F_S - F_B = M \cdot \ddot{x}_N = 0 \tag{5.2}$$

$$J_N \cdot \dot{\omega}_N = -T_S + T_{PH} + T_{RG} + T_B \tag{5.3}$$

The ring gear is only driven when the ball nut rotates in the positive direction, i.e. $\omega_N > 0$.

$$J_{RG} \cdot \dot{\omega}_{RG} = T_{N/RG} - T_{PG} + T_{RS} - T_{Rf} \tag{5.4}$$

For the planetary holder, it is the opposite: it cannot be driven if $\omega_N > 0$. It only rotates in the negative direction.

$$J_{PH} \cdot \omega_{PH} = T_{N/PH} - T_{PG} + T_{PS} - T_{Pf} \tag{5.5}$$

The rotor of the generator is coupled to the sun gear of the planetary gear:

$$J_G \cdot \dot{\omega}_G = T_{SG} - T - T_{Gf} - T_{cog} \tag{5.6}$$

To link the force inside the planetary gear, the forces and moments on the planet gears needs to be calculated:

$$M_{PG} \cdot \frac{(Z_{SG} + Z_{PG})m}{2} \dot{\omega}_{PH} = F_{SG} + F_{RG} - F_{PH} \tag{5.7}$$

$$J_{PG} \cdot \dot{\omega}_{PG} = \frac{Z_{PG} m}{2} (F_{RG} - F_{SG}) \tag{5.8}$$

The indexes used in the equations can be explained by the following list:

S : Ball screw	N : Ball nut	B : Thrust bearing
PH : Planetary holder	RG : Ring gear	PG : Planet gear
SG : Sun gear	f : Friction	cog : Cogging

The forces on the planet gears can be translated to torques:

$$T_{SG} = \frac{Z_{SG} m}{2} F_{SG} \tag{5.9}$$

$$T_{RG} = \frac{Z_{RG} m}{2} F_{RG} \tag{5.10}$$

$$T_{PG} = \frac{(Z_{SG} + Z_{PG})m}{2} F_{PG} \tag{5.11}$$

By substituting in the previous equations, the torque applied onto the ring gear (by the gearing mechanism) can be computed:

$$T_{RG} = \frac{Z_{RG}}{Z_{PG}}\left(-\frac{Z_{PG}}{Z_{SG}}T_{SG} + J_{PG}\cdot\dot{\omega}_{PG/0}\right) \tag{5.12}$$

The torque on the planetary holder then follows:

$$T_{PG} = \left(Z_S + Z_P\right)\left[\frac{1}{Z_R}T_{RG} + \frac{1}{Z_S}T_S - \frac{m^2}{4}\left(Z_{SG} + Z_{PG}\right)M_{PG}\cdot\dot{\omega}_{PH}\right] \tag{5.13}$$

In the ball screw mechanism, the linear motion of the screw is linked to the rotation of the ball nut and can be expressed by the equation:

$$\omega_N = \frac{2\pi}{lead}\cdot\dot{x}_S \tag{5.14}$$

The force required to drive the system is then:

$$F_S = \frac{1}{\eta}\frac{2\pi}{lead}\cdot T_N \tag{5.15}$$

where η is the efficiency.

Similarly, in the planetary gear, the relative velocities of the different components are connected.

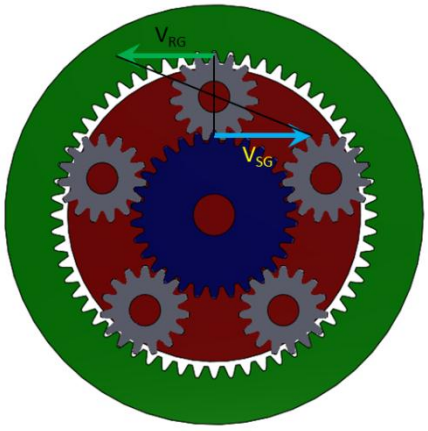

Figure 5.2: Velocities of the planetary gear components

Assuming that the planetary holder is fixed, we have:

$$\begin{aligned} & V_{RG} = -V_{SG} \\ & \Rightarrow r_{RG} \cdot \omega_{RG} = -r_{SG} \cdot \omega_{SG} \quad \Rightarrow \quad Z_{RG} \cdot \omega_{RG} = -Z_{SG} \cdot \omega_{SG} \\ & \qquad\qquad\qquad\qquad\quad \Rightarrow \quad \frac{Z_{RG}}{Z_{SG}} = -\frac{\omega_{SG}}{\omega_{RG}} \end{aligned} \tag{5.16}$$

If the planetary holder is now free to rotate, the previous equation still holds but the speeds are with respect to the planetary holder (not a fixed reference frame), so:

$$\frac{Z_{RG}}{Z_{SG}} = -\frac{\omega_{SG/PH}}{\omega_{RG/PH}} = -\frac{\omega_{SG/0} - \omega_{PH/0}}{\omega_{RG/0} - \omega_{PH/0}} \tag{5.17}$$

This yields the general equation for a planetary gearbox:

$$\frac{Z_{SG}}{Z_{RG}} = \frac{\omega_{PH} - \omega_{RG}}{\omega_{SG} - \omega_{PH}} \tag{5.18}$$

If the sun gear is the output, then it is simpler to rewrite the equation as follows:

$$\omega_{SG} = \left(1 + \frac{Z_{RG}}{Z_{SG}}\right)\omega_{PH} - \frac{Z_{RG}}{Z_{SG}}\omega_{RG} \tag{5.19}$$

where the first half corresponds to the case where the ring gear is fixed, and the second corresponds to the case where the planetary holder is stationary; the general solution is the linear combination of the two.

The clutch bearing, also called a one-way clutch, are devices that allow torque to be transmitted only in one direction, and disengage the driven element from the driving shaft when the latter is rotating more slowly. There are then two distinct modes (Figure 5.3). When the clutch is disengaged, no torque is transmitted. The unknown is then the speed of the driven element. The clutch will re-engage when the speed of the driving shaft is trying to rotate faster than the driven shaft. From that point, the speed of both shafts is identical. The torque is now unknown, and the clutch will then be released when the torque is negative. This system does not permit resistive torque to be applied.

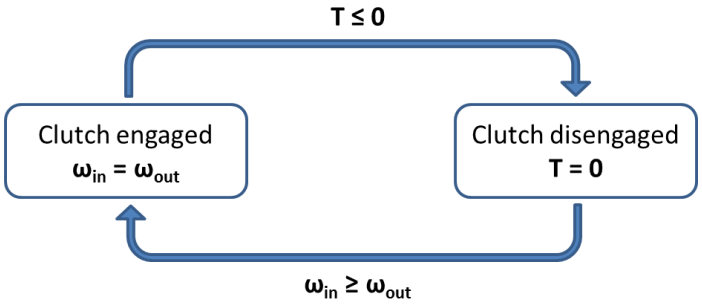

Figure 5.3: Modes of clutch bearings

A model for the generator is also needed. The emf voltage is proportional to the speed of the generator:

$$V_{emf} = k_e \cdot \omega_G = \frac{1}{K_V} \cdot \Omega_G \tag{5.20}$$

where k_e is the back-emf constant, ω_G is the speed of the generator in rad/s, K_V is the motor velocity constant (in RPM per volt), and Ω_G is the speed of the generator in RPM.

The two constants are evidently related: $K_V = \frac{60}{2\pi \cdot k_e}$.

If we assume that the ohmic losses are dominant and that other losses can be neglected, the output voltage of the generator, across a resistor R, can be expressed by:

$$V = \frac{R}{R + R_i}\left(V_{emf} - 2 \cdot V_f\right) \tag{5.21}$$

where R_i is the internal resistance of the generator and V_f is the forward voltage of the diode. It is important to mention that if the electromotive force is lower than the diode threshold voltage, the output is not negative but just zero. The diodes allow current to flow only after the voltage has reached the threshold, usually around 0.5-0.7 volts.

The resistive torque of a permanent magnet generator (or motor) is proportional to the current:

$$T = k_t \cdot I \tag{5.22}$$

where k_t is the torque constant, which is equal to k_e.

So the motor constants are linked by the following relationship:

$$k_t = k_e = \frac{60}{2\pi \cdot K_V} \tag{5.23}$$

With all the equations defined, a model is built in Simulink. One missing element needed to run the simulation is the moment of inertia of each element. Although this is possible to measure, it is neither simple nor easy to do so. The approach used here is to measure the weight of every part with a digital scale, match it in the CAD model, and then use the moment of inertia given by the software.

To test the model, sinusoidal inputs are used, that are identical to those used to test the actual prototype, so it is therefore possible to compare the results. The velocity profiles of the different elements of the planetary gear (Figure 5.4) are as expected. The planetary holder and the ring gear only rotate in one direction. At the end of half oscillations, their speed is higher than the input, proving that the clutches at that point are disengaged. The difference in the gear ratio during rebound and compression is made clear.

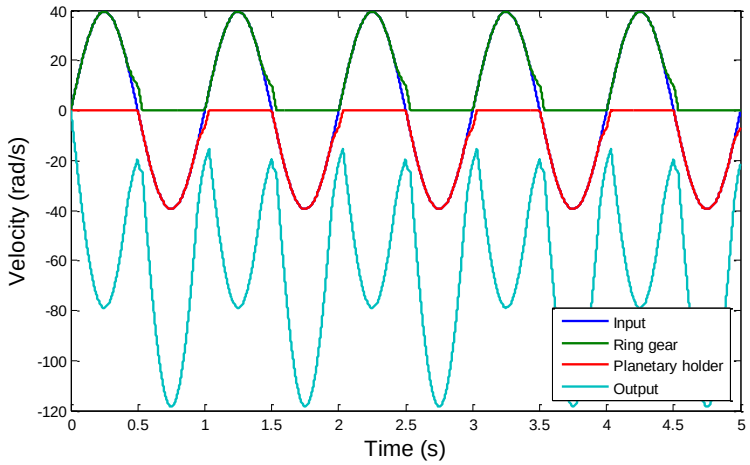

Figure 5.4: Velocity of the planetary gear components

The voltage from the simulation matches the test data relatively well (Figure 5.5). Quite surprisingly, the voltage measured on the actual prototype is higher than the simulation, which may be due to the large variable forces experienced by the prototype. Instead of being gently brought to speed, the generator is quite violently pushed, exceeding the anticipated speed. This can be tracked back to the backlash in the system and to the fact that the parts are not perfectly rigid. Elastic behaviors need to be included to the model to offer a more accurate simulation.

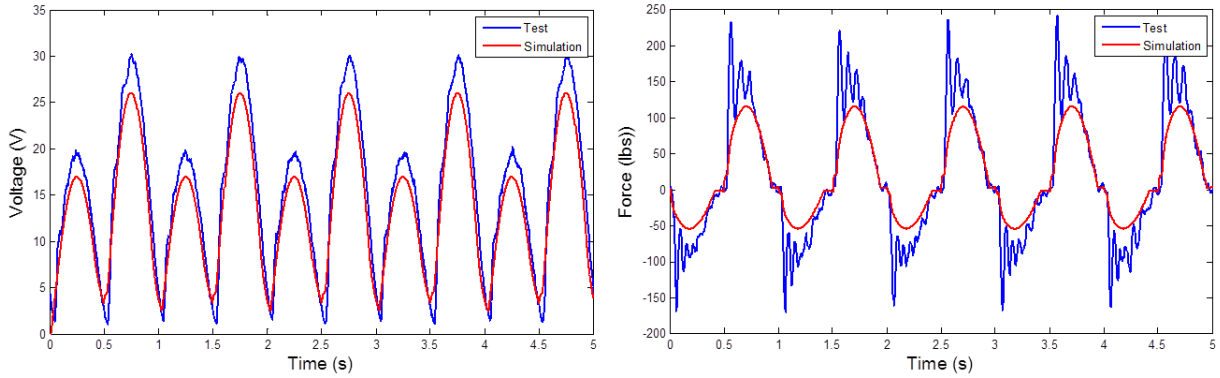

Figure 5.5: Comparison between test and simulation results

The dynamic model is then improved by including an elastic element between the ball screw and the ball nut. The linear motion of the ball screw is translated into a rotation using the same equation as earlier:

$$\omega_i = \frac{2\pi}{lead} \cdot \dot{x}_S \tag{5.24}$$

However, the rotation is no longer directly the rotation of the ball nut, which is now obtained using the following relationship:

$$J_N \cdot \dot{\omega}_N = -T_{Spring} + T_{PH} + T_{RG} + T_B \tag{5.25}$$

where $T_{spring} = k_s\left(\theta_N - \theta_i\right)$

The force on the screw is now given by:

$$F_S = \frac{1}{\eta}\frac{2\pi}{lead} \cdot T_{Spring} \tag{5.26}$$

The force from the bearing includes two components: one constant and one dependent on the position of the nut:

$$T_B = -\left(T \cdot \cos^2\left(a \cdot \theta_N\right) + T_{cst}\right) \cdot sign\left(\omega_N\right) \tag{5.27}$$

The constant component corresponds to dynamic friction, and the variable component corresponds to the defect of parallelism between the bearings, for example, that will generate some hard positions to go through at certain angles. This force is of course opposed to the motion, hence the term $sign\left(\omega_N\right)$ in the expression.

Including those two changes in the Simulink model leads to simulation outputs that better match what has been measured during lab tests (Figure 5.6).

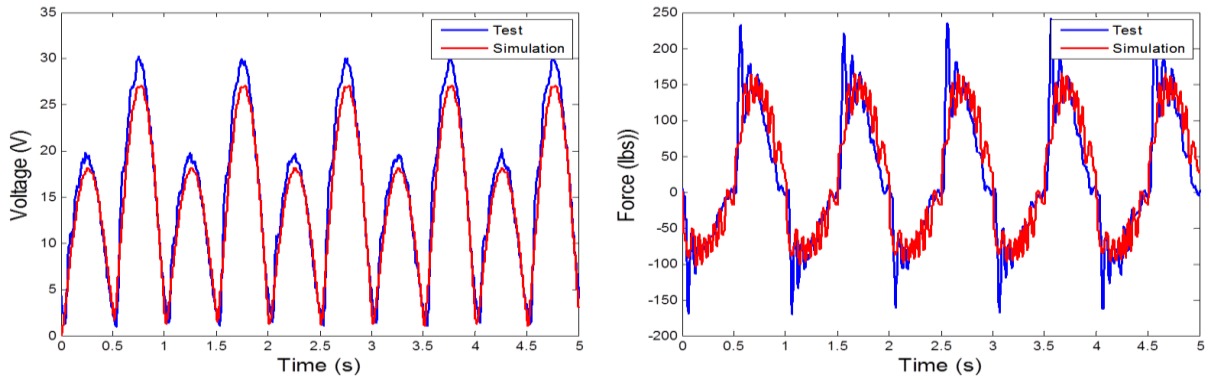

Figure 5.6: Comparison between test and improved simulation results

The waviness matches rather closely the measured forces, although it does not quite perfectly match the first peak. This can be attributed to static friction that has not been included in the model. At each half oscillation, either the ring gear or the planetary holder will start to move again: they usually stop shortly after the ball nut changes direction. The force required to get them started can explain the large peak at every half-oscillation. More sub-functions could be added to the model but, instead of setting the goal of perfectly matching the model to the data results, efforts should be made to bring the behavior of the prototypes closer to the ideal case. The implementation of rubber elements on both sides of the ball nut bearings is an example of that process. It is worth mentioning that two prototypes will never give the same wavy pattern, and even one given prototype can change during testing. All contribute to making a universal model for the harvesters difficult to achieve. However, the real goal of the simulation is to accurately estimate both the expected voltage/power generated and the forces required. In that respect, the simulation performs rather well.

The match of the output voltages also improves compared to the first simulation. It is believed that modeling the large peaks seen in the forces would close the gap even more.

Similar resisting torques can be added to all the rotating parts. In simulations, they lead to similar results, adding wavy pattern to the output voltages and forces.

A low pass filter can be added to simulate the effect of the capacitor connected after the diode bridge, in parallel with the resistor.

Chapter 6

Miniature Energy Harvester

6.1 Introduction

The current prototype designed to fit inside a freight car spring is derived from the same concept as the first generation: using a ball screw to transform the linear motion of the car suspension into rotation, and then using a special gearing arrangement to keep the generator always rotating in the same direction. With room available inside the D5 spring, the spring commonly found in most freight car suspensions in North America, it was possible to scale the prototype up in size. The prototype could be further increased in size to meet the requirements of a particular application. Generally, more room gives more options in the choice for the different system elements: ball screw, one-way bearings, planetary gear, generator, etc.

In numerous other applications, a smaller package might be more desirable. Scaling the concept down might then present some challenges, especially when it comes to finding standard elements that can fit the applications. This chapter will focus on the development of an energy harvesting device using the same concepts, but specifically oriented to limited available space.

The chosen application is retrofitting an existing mountain bike fork. In addition to providing the desired damping, the fork will be able to generate some electrical power that can be used to power the bike computer, LED lights, a GPS, a phone, or any other devices that might be useful on a mountain bike.

6.2 Energy Harvester for a Mountain Bike

6.2.1 Introduction

Mountain bikes are bicycles specially created for off-road cycling, which includes riding on dirt tracks, forest trails and unpaved roads. With the special conditions encountered off-road such as traversing rocks or roots and logs, these bikes are significantly different from their road counter-parts. To withstand the higher stresses that will be faced, they are beefier, with larger frame tubes, and thus heavier. They are also

equipped with large knobby tires that will increase traction and comfort on uneven terrains, with the trade-off of higher rolling resistance. To improve climbing capabilities without sacrificing downhill speed, mountain bikes have been fitted with gearing of up to 27 or 30 speeds: 3 chainrings on the crank, and 9 or 10 sprockets in the rear. All the speeds are not really usable, as they would put too much stress on the chain. Therefore, the current trend is now to reduce the number of speeds to 20 with two chainrings, or even to 10 with only one chainring eliminating the need for a front derailleur.

Since the mid-1990s, front wheel suspensions have become more and more common. Derived from motorcycle technologies, they improve comfort by insulating the rider from the terrain, but more importantly, they improve handling by keeping the wheels in contact with the ground. The characteristics depend on the type of biking for which the fork is designed. Nowadays, forks for cross-country riding usually have 80-100mm of travel. Downhill forks must handle the most extreme conditions with big jumps and rough terrain. So, it is not rare to find forks offering travel of about 200mm. Some forks offer technical advances like travel adjustment, which can be useful when traveling in uphill and downhill sections (particularly common in cross-country), and in lock-out systems which eliminate the fork's motion leading to better efficiency on smooth terrain.

Suspended forks consist of two main parts: a spring and a damper. The spring can be a coil spring made out of steel or titanium, an elastomer, or an air spring. The elastomer element gives, in addition to stiffness, a little bit of damping. On small travel suspensions usually found on hybrid bikes, they can be used alone, but on mountain bikes, they are used in series with a helical spring. Air springs lead to lighter forks but also results in an increased complexity, which translates to a higher price tag. Changing the pressure allows easy adjustment for the rider's weight. With coil springs, a preload adjuster is typically used.

The damping system works by forcing oil through small orifices and shim stacks, thus dissipating energy. Without such a system, the suspension would rebound excessively and the handling would suffer. It is generally possible to adjust the compression and rebound force using knobs at the top and bottom of the fork legs.

Today, it is fairly common to find full-suspension mountain bikes. In addition to the suspended front fork, the rear wheel is also suspended. This is extremely important for downhill riding, but often cross-country riders prefer to sacrifice comfort for efficiency, as the rear suspension will dissipate some of the energy of the pedaling in what is called pedaling-induced bobbing. Full-suspension bikes are also heavier, more complex, and more expensive.

This chapter presents an energy harvester system that fits inside a suspended front fork where the damping would normally be located. A similar idea could very well be developed for the rear suspension.

6.2.2 Presentation of the Mountain Bike Fork: Manitou SX Ti

Previous projects at the Center for Vehicle Systems and Safety (CVeSS) have used Manitou's SX forks (Figure 6.1) to develop semi-active suspensions for mountain bikes using magneto-rheological (MR) fluid technology. With this particular model of suspended fork readily available in our laboratory, the choice has been made to use it as a base for a retrofit with an energy harvesting system.

Figure 6.1: Manitou SX Ti

As can be seen on the fork's schematics (Figure 6.2), the right leg houses the suspension system composed of a coil spring and an elastomer in series, while the left leg houses the damping system. This particular fork uses the Twin Piston Cartridge (TPC) system. One piston is used for compression

damping, and the other is used for rebound. Both create damping by forcing oil through a shim stack, and a by-pass can be adjusted externally by the rider to adapt to the terrain conditions. This system is still in use in today's forks of this brand.

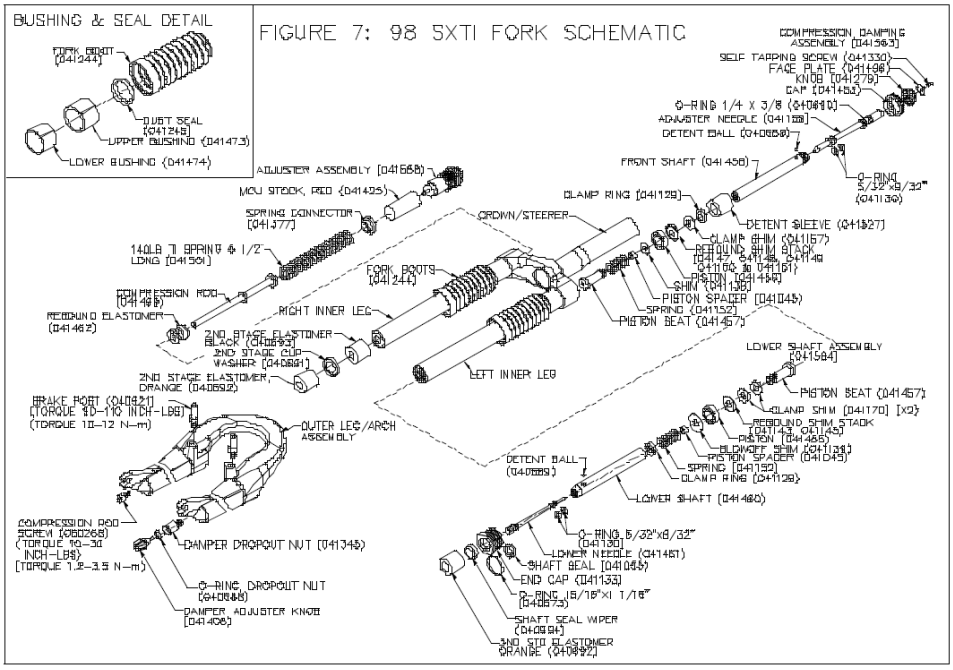

Figure 6.2: 1998 Manitou SX Ti schematic [50]

With the spring assembly and the damper well separated, the SX Ti is a perfect candidate to be retrofitted with an energy harvesting system. The right leg with the spring-elastomer arrangement will remain untouched, while the damping system will be completely removed, leaving the necessary space for the harvester.

This particular fork is rather old (about 15 years), but the design of today's forks remains similar: one leg for the spring arrangement, and one leg for the damping. This can be seen on the schematic of the 2012 Manitou Circus Expert (Figure 6.3). It can be noted that the damping system is now in the right leg. The main difference between the 1998 fork and today's fork is the travel. In 1999, the Manitou SX travel increased from 70mm to 80mm. Today, the Manitou Circus is offered in 100mm (internally adjustable to 80mm) and in 130mm for the *Expert* model. This is typical of the travels found on forks for cross-country riding. For increased stiffness, the stanchions have grown from 28mm (outside diameter) to 32mm. Therefore, if it's possible to fit a harvesting system to the 1998 fork, it should be no issue fitting it into the current model.

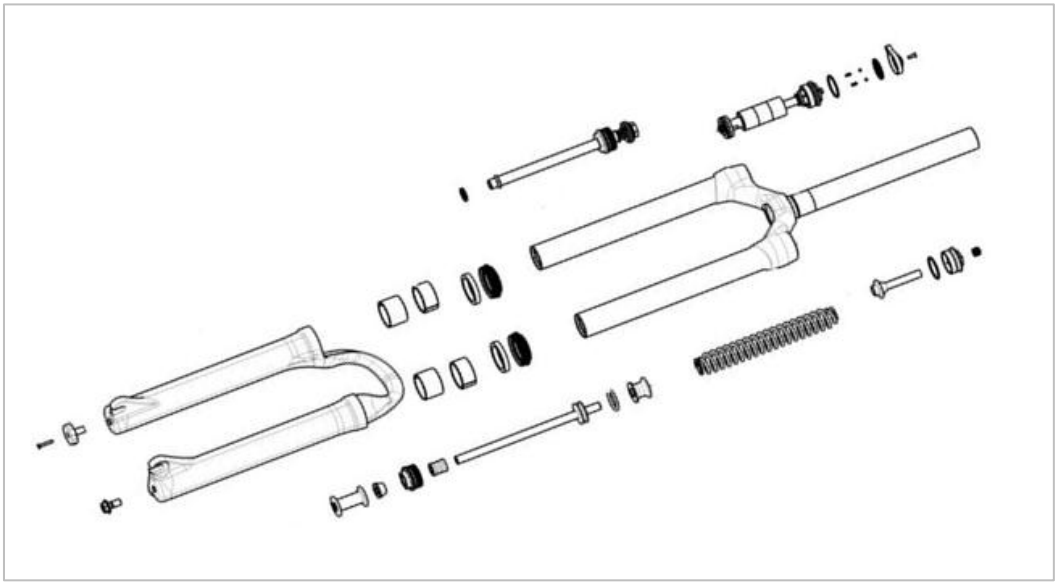

Figure 6.3: Circus Expert fork schematic [51]

6.2.3 Energy Harvesting System

The harvesting system developed for the mountain bike application reuses the proven concept developed for freight car suspensions. As presented before, the ball screw converts the linear motion of the suspension into a rotation, and a smart gearing arrangement using one-way bearings then transmits the motion to the generator so that it will always turn in the same direction.

Esthetically from the outside, the fork looks completely unchanged. But, as mentioned earlier, the energy harvesting system is housed in the left leg in place of the damper (Figure 6.4). The spring arrangement remains in the right leg, exactly how it is received from the factory.

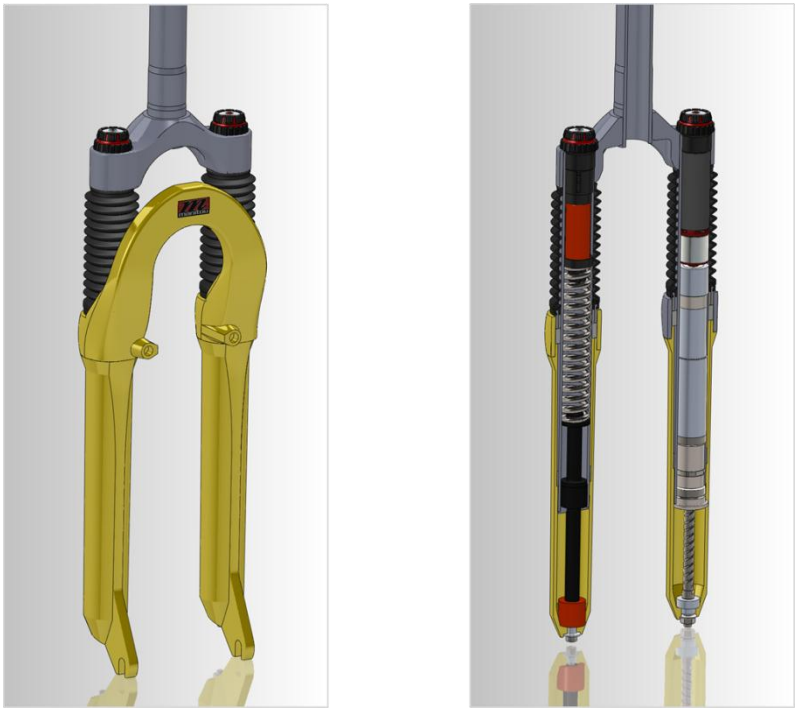

Figure 6.4: CAD model of the fork: complete and cut-away

Since the same concept is used here as in the freight car suspension, the same elements can be found in both prototypes (Figure 6.5). With the generator placed at the very top, the design can actually be considered closer to the prototype of the first generation.

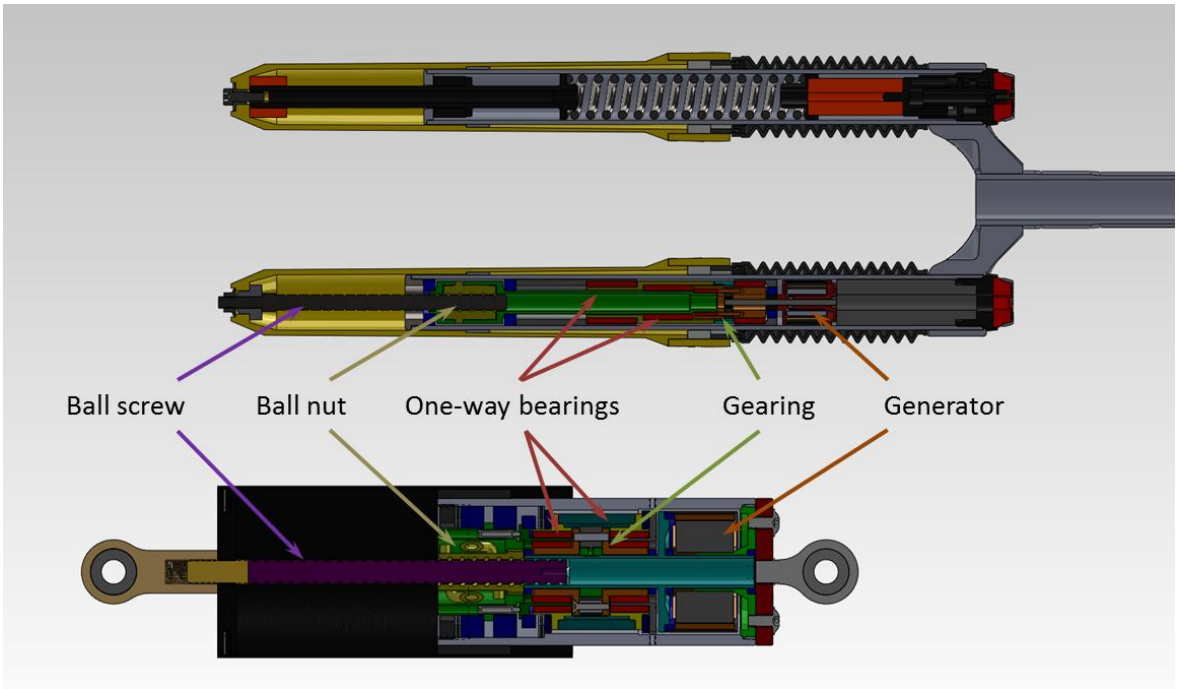

Figure 6.5: Side-by-side comparison of the energy harvesters

Although this particular fork has only 70mm (2.75 inches) of travel, the energy harvesting system is designed to have a stroke of 100mm (4 inches), allowing the system to be directly fitted in more modern forks. If even longer travel is required, this can easily be achieved as there is a 50mm spacer on top of the generator. Besides being used to increase travel, 50mm (2 inches) also give the options of accommodating a longer generator and adding an additional gearbox stage to further increase the generator speed.

The complex issue in adapting the concept to a different space requirement is mainly in finding parts that can fit the application. The current prototype for railroad applications is designed to be placed inside a 3.5-inch ID spring. To fit in the leg of the fork, the new prototype cannot be larger than 24mm (0.945 inch).

It was again planned to use a repurposed DC brushless motor for the generator. The RC hobby industry provides numerous options in the size of these efficient motors (Figure 6.6): their outside diameters range from 13mm (0.5 inch) and smaller to more than 100mm (4 inches).

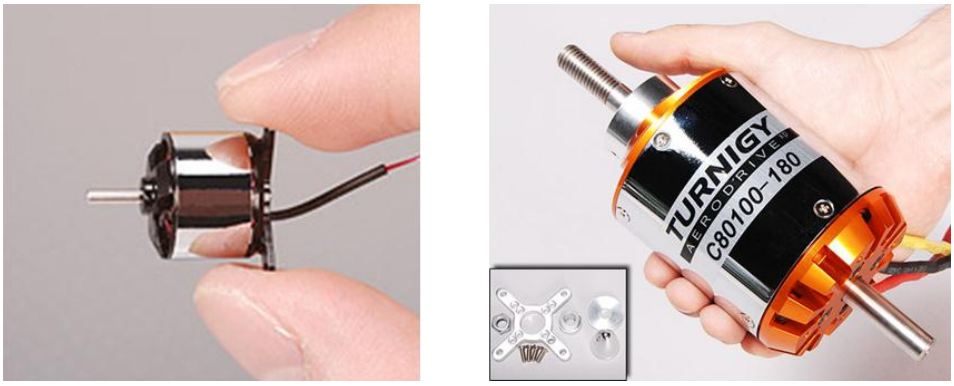

Figure 6.6: Examples of extreme size RC brushless motors [52]

For this particular application, the ADH 300XL was chosen. With an outside diameter of 23mm, it is the perfect size to fit in the fork leg and maximize the room available. With a relatively high Kv value, it may need to be rewound to fit the application. As a stock motor, it has a nominal voltage of 11.1V (3 cells Li-Po battery), and can draw up to 44 amps to reach around 450 Watts.

The one-way bearings make the unidirectional motion of the generator possible. In the railroad prototypes, bearings of inside diameters of 25mm (0.98 inch) and 35mm (1.38 inch) are used. They are obviously too big to fit in the fork, so new sizes are needed. One-way bearings are available in 11 metric sizes, ranging from 6mm to 35mm (inside diameter) and in 7 inch sizes, ranging from 1/8 inch to 1 inch. It is then possible to obtain available standard sizes of one-way bearings that can fit in the fork leg.

The planetary gearbox is another essential element as this is where the motion becomes unidirectional. Unlike one-way bearings, gearboxes are not available in every size, but some companies, like *GYSIN* distributed by *Precision Planetary Gears* [53], offer options that would be perfect for this application. Their planetary gearbox GPL-22 (Figure 6.7) has an outside diameter of 22mm, maximizing the 24mm available inside the fork leg.

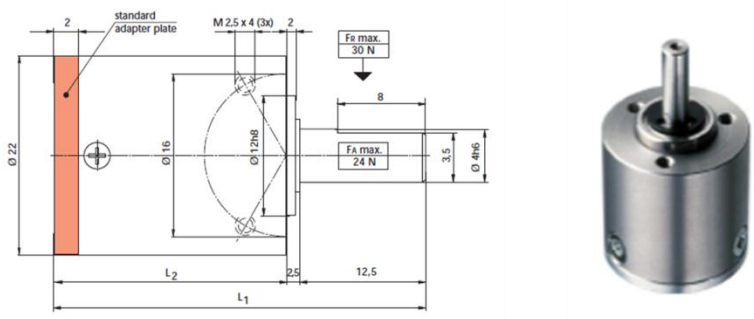

Figure 6.7: GYSIN planetary gearbox GPL-22 [53]

The last main element of the prototype is the ball screw/ball nut. The Japanese company KSS [54], distributed in the U.S. by *The Precision Alliance* [55], for example, has an enormous selection of ball screws. They are offered in numerous diameters and various leads. For the bike application, a ball screw with an 8mm diameter and a 12mm lead has been chosen. A larger diameter allows for higher load capacity. At 8mm, the screw is rated for 4000N (900lbf), which should be more than enough for the application. A larger lead limits the stress on the screw, which is here backdriven. At 12mm, it is quite close to the 0.5-inch used in the current design.

Various bearings are also required, but they come in almost every possible size and are relatively easy to source.

The ability to find standard parts that meet the requirements of a particular application is always very important. This usually reduces the design/study phase, the time to build a prototype, and the cost. It also makes maintenance and repairs easier, as parts can be ordered if they need to be replaced. Using custom-made parts is usually more expensive and requires an added phase of development and reliability testing. The time required to get the parts can be an issue, especially after failure. Finally, it is sometimes nearly impossible to make custom parts. For example, it would be extremely difficult to make custom-size one-way bearings. Therefore, it was really important in the design process of this miniaturized prototype to find off-the-shelf parts.

6.2.4 Conclusion

Through this study, it has been proven that it is possible to design a small-size energy harvesting system based on the same concept used in the well-proven prototype for railroad applications. Mostly off-the-shelf elements were used, and custom-made parts were limited to a minimum. Designed to fit in the 1998 Manitou SX fork that was available in our laboratory, the system can be fitted to virtually any mountain bike fork, given that it has one leg dedicated to the damping system.

The dynamic model built for the freight car suspension system can be directly used to simulate the behavior of this prototype. The only changes would be in the Matlab files that contains the masses and inertias of the various parts. With those adjustments, the simulation would run the same way, without any changes required.

For a successful prototype, it is necessary to conduct a proper campaign of simulations to determine the required characteristics of the generator and the gearing system, since the conditions that are encountered by a mountain bike are radically different from those in a freight train. The input velocities should be significantly higher, but the forces available would be immensely reduced.

6.3 Concluding Remarks on Miniature Harvesters

Fitting an energy harvester into a mountain bike fork is a concrete demonstration that it is possible to adapt the concept to applications with limited space. This is just one possible example, and many more applications could be considered. The versatility of the design makes harvesting energy possible in virtually any situation where linear motion is present.

This page intentionally left blank.

Chapter 7

Rack and Pinion Design

7.1 Introduction

The work presented in the previous chapters has focused on developing energy harvesters that use a ball screw system to transform the linear motion of the suspension into a rotation. As mentioned at the beginning of Chapter 4, the very first attempt at using a rotating generator (Figure 7.1) was with a rack and pinion [56].

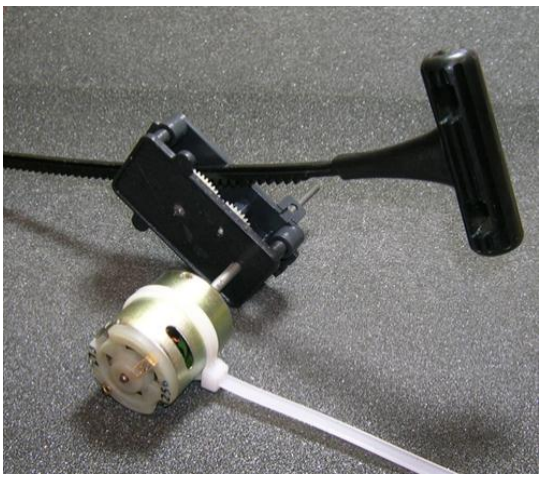

Figure 7.1: First attempt at using a rotating generator with a rack and pinion

At that time, in the spring of 2009, various concepts were developed around a rack and pinion. The ball screw was then pursued, as it allows a more elegant design with a packaging similar to a regular shock absorber, and it also leads to a denser system with a larger generator. Systems using racks and pinions have their own merits, and this chapter will show why they are worth being investigated by presenting the different concepts imagined and the resulting complete design of a prototype that can be fitted inside a freight car suspension.

7.2 Unidirectional Motion of the Generator

7.2.1 Introduction

In previous chapters, the benefits of having the generator always rotating in the same direction have been discussed. Mechanically, this leads to lower forces by keeping the inertia of some of the rotating parts. Electrically, it yields a smoother output. A system to allow this function has been implemented in the ball screw prototype, and it is imperative that it be included in any rack and pinion design.

7.2.2 Concepts

Unidirectionality of the generator is rather complex in the ball screw prototype because the rotations of the different moving parts are all concentric. With shafts rotating on parallel axes, such as it is the case when using racks and pinions, unidirectionality becomes easier to achieve.

Like the previous concept, one-way bearings will be used to transmit rotation in only the right direction; these are represented in red in the illustrations.

The first concept is to have two pinions on one single rack (Figure 7.2). The two pinions are driven back and forth by the rack. On each shaft of the pinions, a gear is placed and can only be driven in one direction through a one-way bearing. Therefore, when the rack is moving downward, both pinions rotate counter-clockwise, but only one gear (in blue) on a pinion shaft is driven: the bottom gear in Figure 7.2. The light blue gear is then rotating clockwise. Conversely, when the rack is moving upward, the pinions rotate clockwise; the blue gear on top is driven through the one-way bearing, and the light blue gear is again rotating in the clockwise direction. In this case, the motion is transmitted through the bottom blue gear. Both blue gears are always rotating in the same direction as well: the top one is always rotating clockwise, and the bottom one is always rotating counter-clockwise. However, they are only transmitting the power coming from the rack through the one-way bearing only in one direction.

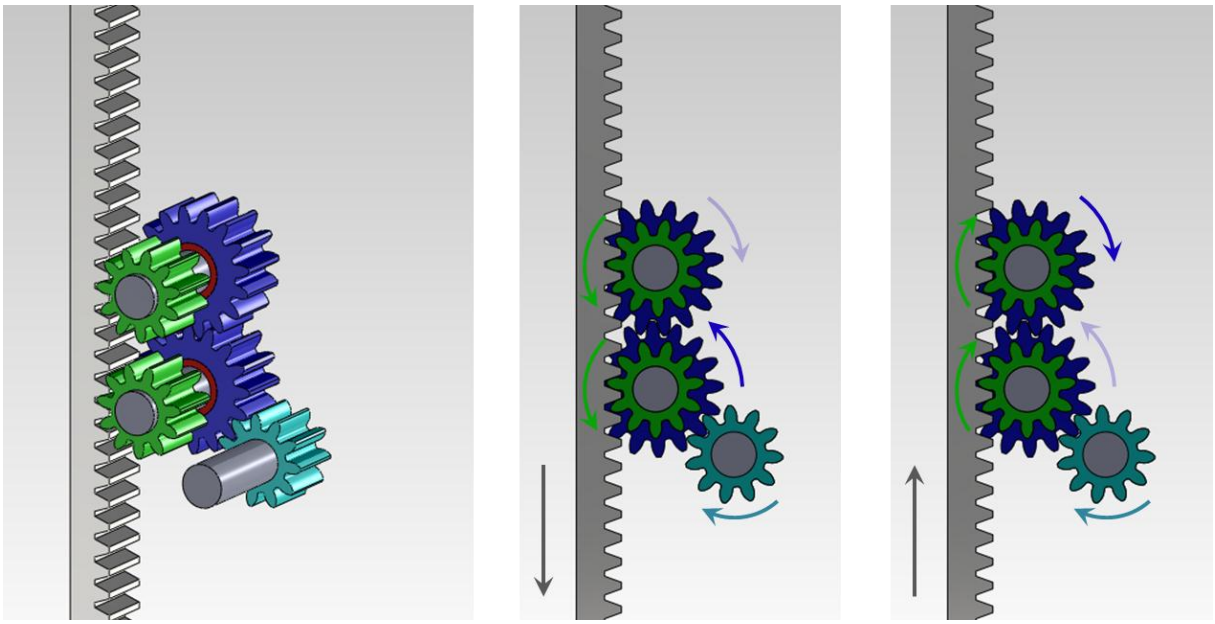

Figure 7.2: One rack - two pinions concept

Obviously, the clockwise rotation could be transmitted through the bottom (blue) gear and the counter-clockwise rotation by the other one, on top. This is totally arbitrary: the only difference it would make is that the light blue pinion would always rotate counter-clockwise. One of the two cases may be more desirable depending on how the generator is connected.

In Figure 7.2, the one-way bearings are housed in the gears (in blue). They could also be placed inside the pinions, if their size allows. The pinions are important in the overall gear ratio of the system, and in order to have faster rotations, they need to be as small as possible. This will, in most cases, force the one-way bearings to be in the gears rather than the pinions.

The drawings are for illustration purposes only. It is clear that the bearings and supports that would be needed are not shown here.

A second concept is to use two racks and two pinions (Figure 7.3). The two pinions are driven back and forth by two different racks, in contact on opposite sides. During a downward motion, the left pinion is driven counter-clockwise, whereas the right pinion is driven in the opposite direction. In upward motion, the directions of rotation are inverted. Again, the blue gears on the pinion shaft can only be driven in one direction because of one-way bearings. So the power from the rack is transmitted through the left blue gear when going downward, and through the right gear during upward motion. The result is that both blue

gears are always rotating in the counter-clockwise direction, and the light blue gear is always rotating clockwise.

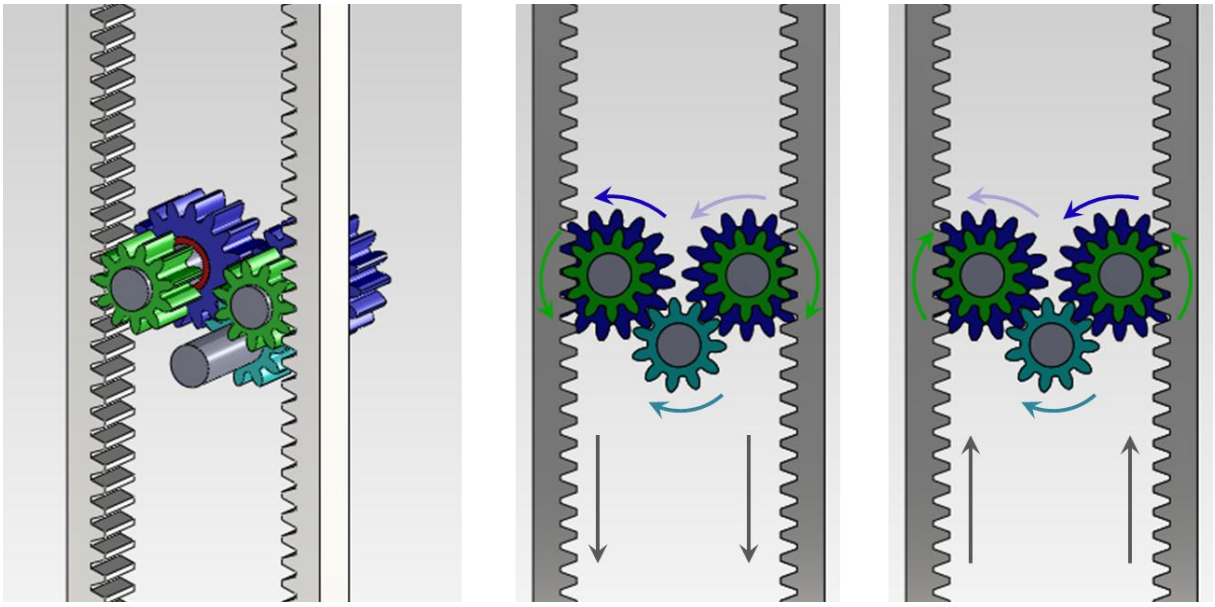

Figure 7.3: Two racks - two pinions concept

Just as in the previous design, the directions of rotation could easily be inverted if needed.

In the example in Figure 7.3, the blue gears are on the same side of the rack and pinions. In some designs, it may be more convenient to have one on each side, which would limit interferences and allow larger sizes of gears.

Another variation would be to have the pinions (in green) in contact with each other. Thus, the force transmitted through the one-way bearings would come from both racks and both pinions. This would allow twice the force to be applied or, for a given force, this would decrease the stresses on the pinions by a factor of two. Another consequence is that since the force is then shared by both racks, the overall force would be centered, putting the mounting points in perfect alignment: no more bending moments on the racks, and no alternate side-to-side loading on the system.

In Figure 7.3, the pinions (in green) are on different axles; however, it would be possible for them to share that axle (Figure 7.4). Only one gear would then be needed to carry out the unidirectional motion. This requires the one-way bearings to be inside the pinions.

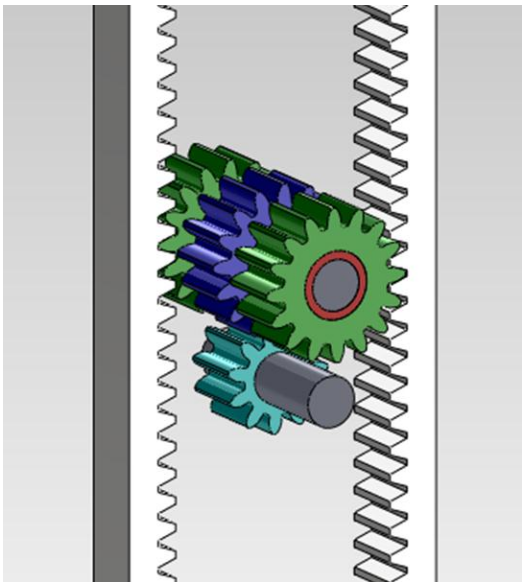

Figure 7.4: Two rack - two pinions on the same axle

A third option is to have only one pinion and one rack. The pinion is driven back and forth by the rack. Two one-way bearings on the axle of the pinion transmit that motion: one during clockwise motion, and the other the rack. The arrangement would be similar to the first axle in Figure 7.4, except that the rack would be driving the blue gear, and the green gears would be the elements always rotating in the same direction. Then the motion would be transmitted to the generator if we ensure that on one of the two paths from the one-way bearings, there is an extra eternal contact to allow both motions in the same direction.

7.2.3 Competing Work

Dr. Nelson and his team at Nebraska-Lincoln have worked on a system aimed at harvesting energy from the vertical motion of a railroad track when a train is passing. They presented their first prototype at the SPIE conference in San Diego, CA in 2009 [57]. Fitted with only one one-way bearing, it was only harvesting energy during downward motion. In 2011 [58], they presented their improved prototype (Figure 7.5) which, using two one-way bearings, is then able to harvest energy during downward and upward motion. The system uses a one rack/one pinion design. The pinion shaft is directly connected to a one-way bearing, and is linked to a second bearing through a set of gears. This leads to a motion of the gearbox and thus the generator always in the same direction.

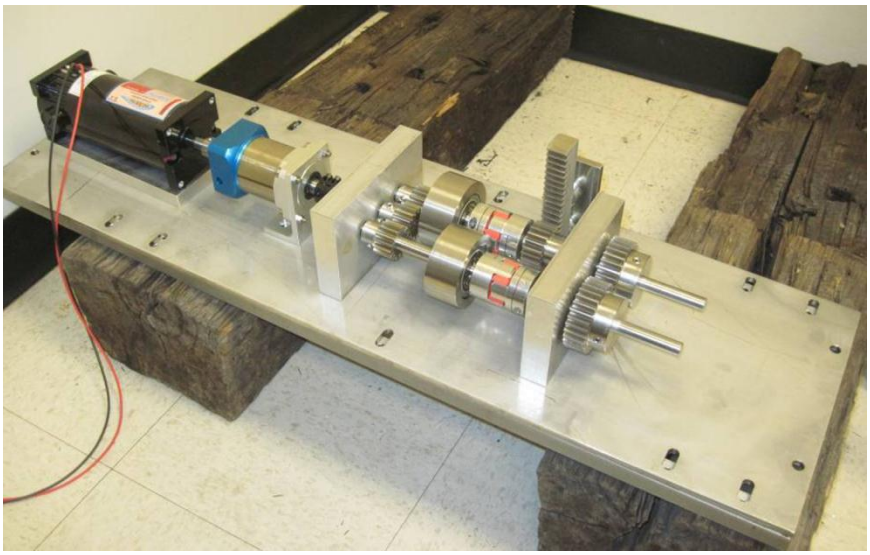

Figure 7.5: Railroad energy harvester for wayside applications from Nebraska - Lincoln

Dr. Zuo and his team at Stony Brook University have been working on energy harvesting systems for car suspensions and for vertical track deflections. Their energy-harvesting "shock absorber" (Figure 7.6) uses a rack and pinion system to transform the linear motion of the suspension into rotation. A smart arrangement of bevel gears then converts the alternative motion of the pinion in a unidirectional rotation. By analogy with a diode bridge, they call the system a mechanical motion rectifier (MMR) [33].

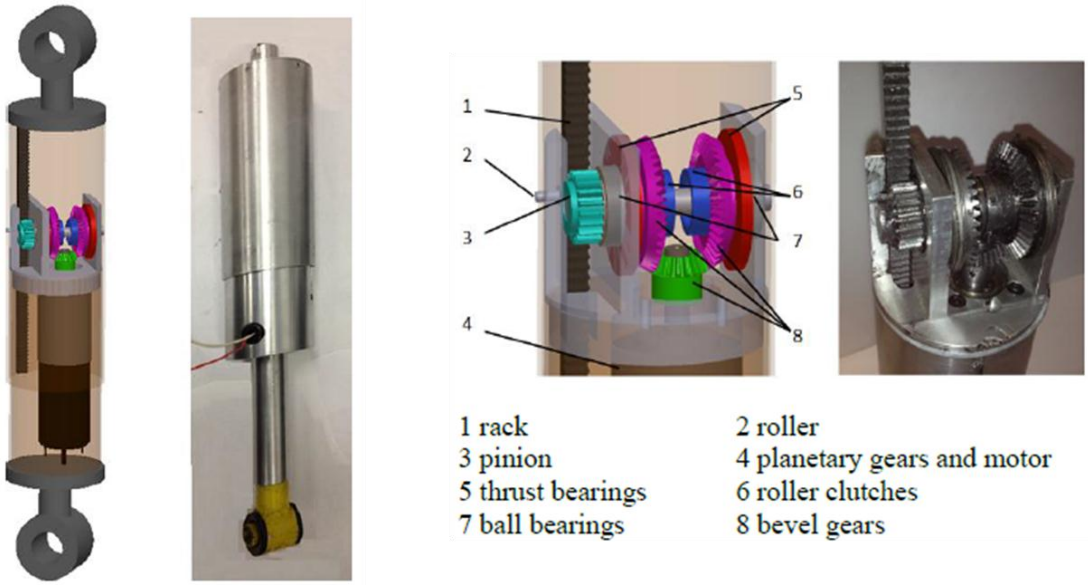

Figure 7.6: Energy-harvesting "shock absorber" from Stony Brook [33]

The bevel gears are driven through one-way bearings. In one direction, one of the two purple gears is driven. Conversely, in the opposite direction, the other transmits the torque. The result is that the green gear always rotates in the same direction, and thus so does the generator. This prototype seems to be giving promising results in laboratory and field tests.

The same team has also been working on a system to harvest energy from the vertical deflection of a rail track (Figure 7.7) [59]. Similar to Nelson's concept, it varies on some points. It uses two rack and pinions instead of just one, and some of the design flaws are corrected and efficiency is improved by adding a large flywheel.

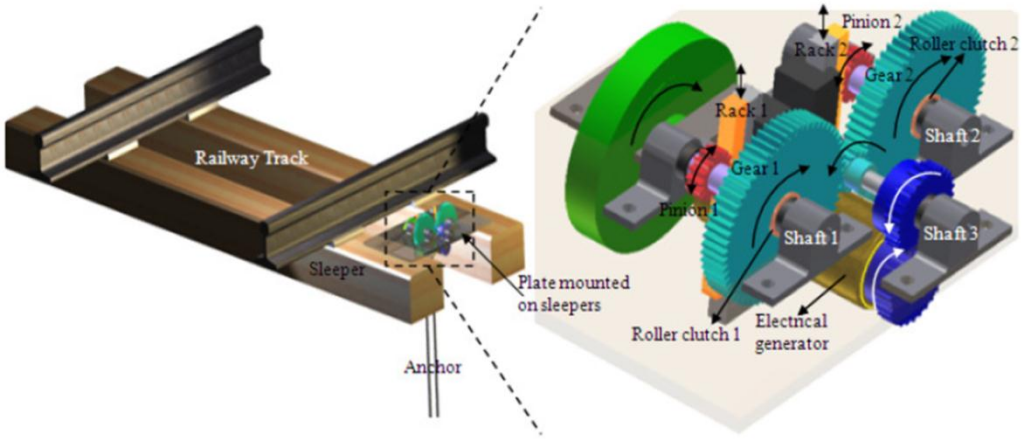

Figure 7.7: System harvesting energy from railroad track vertical motion - Stony Brook [59]

For this concept, Dr. Zuo and his team won the award for "Best Application of Energy Harvesting" at the 2012 Energy Harvesting & Storage Conference, held in Washington, DC on November 7-8, 2012. Those two designs are now protected by a patent, published on January 10, 2013 [34].

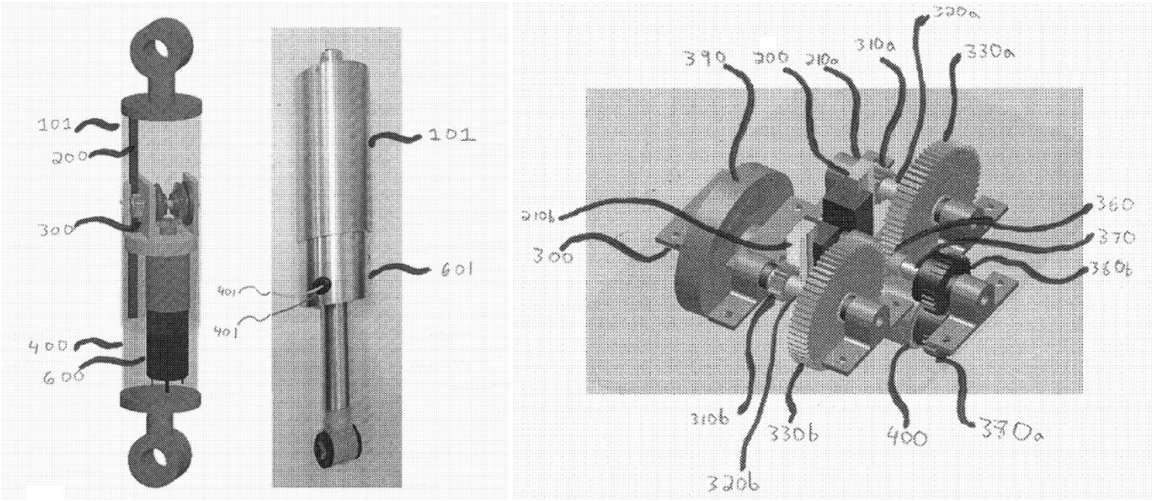

Figure 7.8: Patent US 2013/0008157 filed by Zuo et al.

This patent attempts to be as broad as possible, both in the possible applications (harvesting energy from rail track motion, from buoy, or from flaps) and in the rack and pinion configurations. Most of the concepts presented before are now covered in this patent (Figure 7.9). This will affect the design of the energy harvester for freight car suspension and will be discussed in detail later in this chapter.

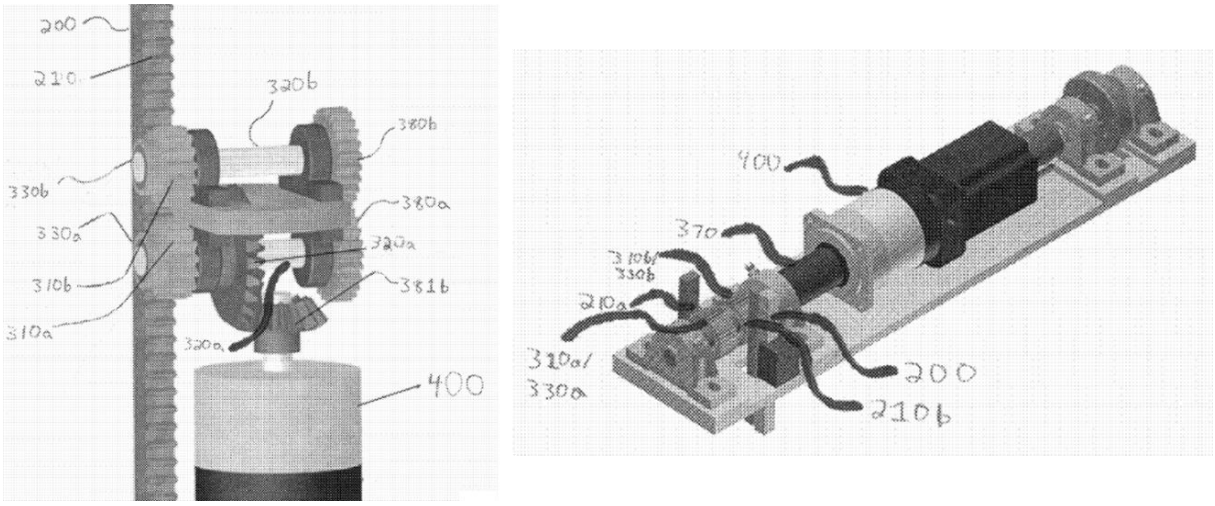

Figure 7.9: Figures 8 and 10 from Patent US 2013/0008157 filed by Zuo et al.

Some one-way bearings manufacturers like Timken and NSK have some interesting examples in their catalog (Figure 7.10) that can be useful in the design of energy harvesting systems.

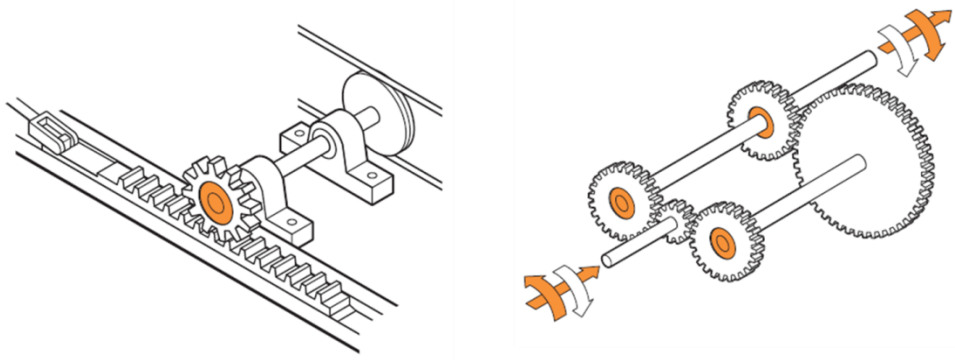

Figure 7.10: Application examples from Timken's roller clutches catalog [60]

7.3 Energy Harvester for Railroad Applications

7.3.1 Introduction

In this section, the concept of an energy harvesting system using a rack and pinion will be presented. The goal was to design, as an alternative to the system presented in Chapter 4, a harvester that is still using the motion of a freight car suspension as input, and that fits inside a D5 spring as well. The rack and pinion system offers some advantages over a ball screw, but it also comes with a new set of design constraints.

7.3.2 Initial Prototype

The ball screw creates a rotation that is directly around the main axis of the system. This has several advantages: first, that the ball screw can easily (and should be) centered, and second that the generator can be as big as the outside tube, maximizing space and power generation potential. The system in Chapter 4 represents this idea quite well. On the other hand, the rack and pinion combination will generate a rotation perpendicular to the main axis of the system. This means that it will be rather complex to have the rack centered, and that if one wants to have the generator oriented in the same direction as the harvester, bevel gears will be needed to change the direction of the rotation. For long travels, the generator will have to be off-centered in order for the rack to have enough room. This limits the possibilities to smaller size generators. The energy-harvesting "shock absorber" from the team of Stony Brook University is a good example of that design.

The idea for the energy harvester presented in this chapter is to keep all the rotations on parallel axis and not to use bevel gears. This means that the generator will be rotating around an axis perpendicular to the main axis of the system (Figure 7.11). In this orientation, the generator cannot be extremely large, but in a 3-inch diameter tube (similar to that of the ball screw prototype and small enough to be inside a D5 spring) it is possible to quite easily fit a generator like the Scorpion S3016 that has been used in the first prototype with a rotating generator presented in Chapter 4.

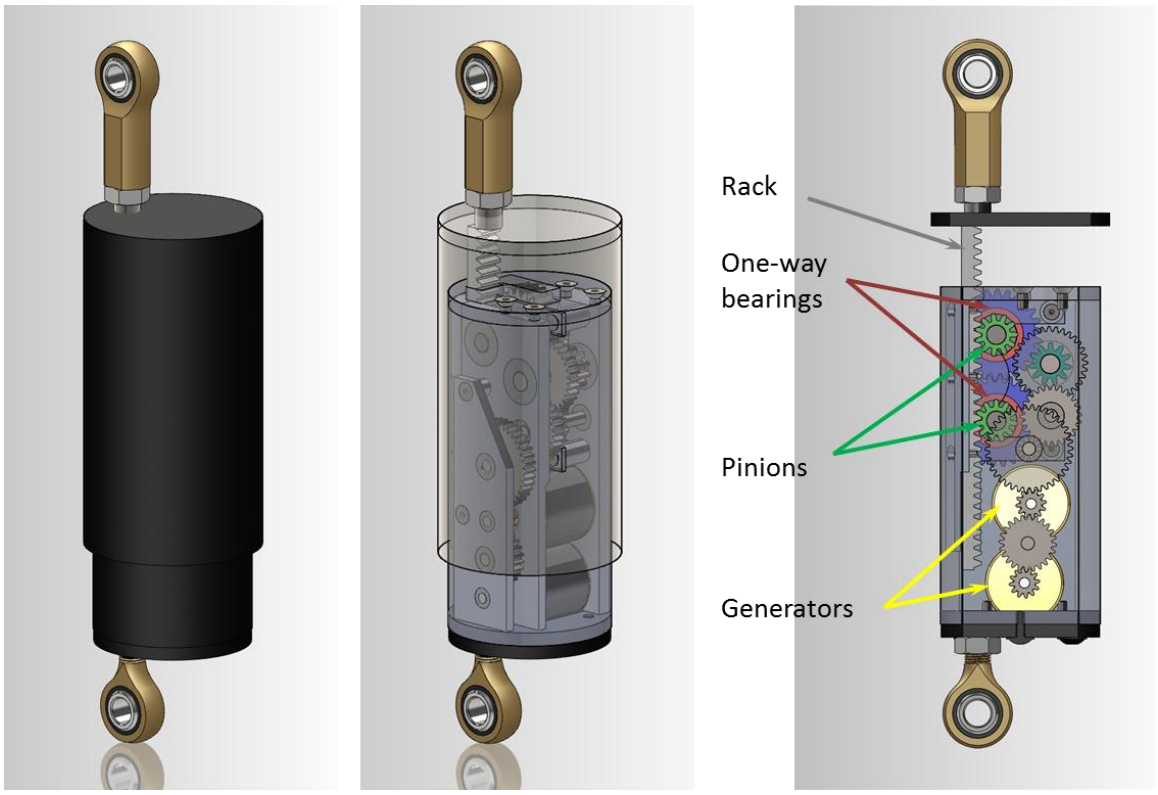

Figure 7.11: Rack-and-pinion energy harvester for freight car suspension

To compensate for the smaller size of the generators, two of them are used in this application. They need to be synched in order to have collaborative effects, in the same manner as generators in power stations when they are connected to the power grid.

In Figure 7.11, it can be noted that the mounting ball joints are off-centered. They are placed that way so they align with the contact force developed between the rack and the pinions. This way, there is no bending moment on the rack or on the caps. The ball joints are primarily integrated in the design for laboratory testing. Inside the spring, another mounting system with welded caps is envisioned.

This prototype is developed around the concept of one rack driving two pinions, presented at the beginning of this chapter. The two green pinions (Figure 7.12) are driven back and forth, and one or the other of the two blue gears is driven depending on the direction of motion of the rack. The result is that the light blue gear is always rotating in the same direction. The speed is then increased through a set of gears before reaching the generators.

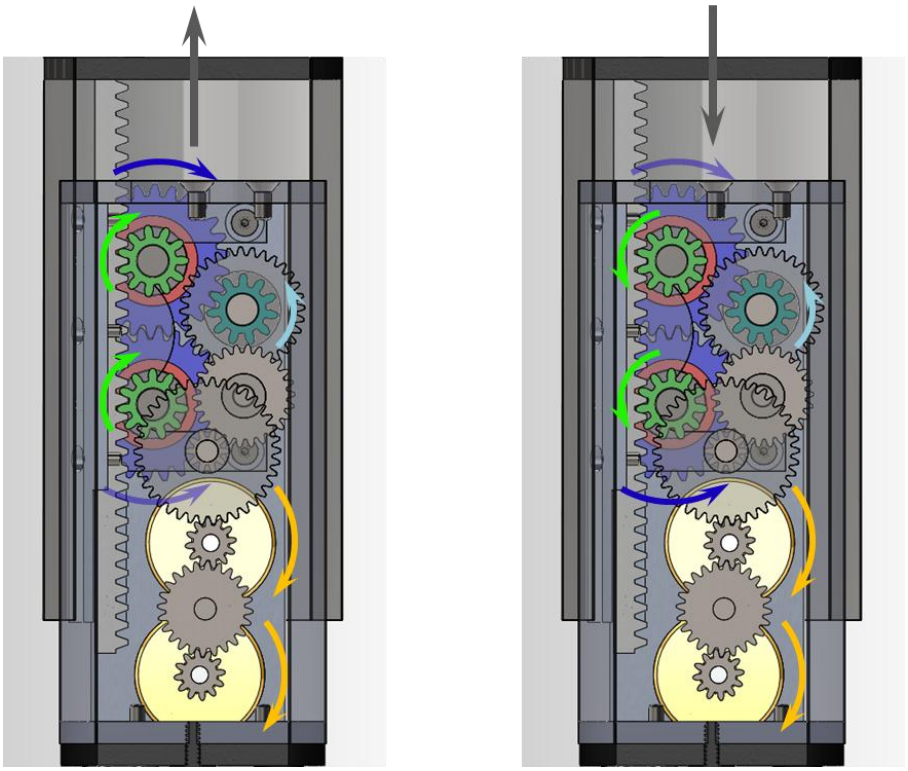

Figure 7.12: Illustration of the motions of the harvester gearing system

In order to limit the gear ratio, the pinions must be as small as possible. The ball screw used in the current prototype has a lead of 0.5 inch, so the ball nut does two full rotations per inch travelled by the screw. To obtain the same values with a rack and pinion, the diameter of the pinion needs to be:

$$R_{pinion} = \frac{lead}{2\pi} \approx 0.080\text{in} \tag{7.1}$$

It is highly unlikely to find a gear with a pitch diameter of 0.160 inch, especially to carry the loads that are expected in freight car suspensions. Larger gear teeth can evidently withstand higher loads, so in that respect, a smaller pitch (or higher module) is desirable. A good compromise for this application is a pitch of 16. In that size, the smallest number of teeth readily available is 11. This yields a pitch diameter of:

$$d_p = \frac{N_{teeth}}{p} = \frac{11}{16} = 0.6875\text{in} \tag{7.2}$$

This means that the rack needs to travel 2.16 inch to obtain one full rotation of the pinion, over four times more than with a ball screw. Besides the kinematic aspect of the gear sizing, the gear has to be chosen so

that it can withstand the loads. The ball screw is rated at 750-1000lbf, depending on the manufacturer. The target for the gear is then set at 750lbf. The stress is then estimated using the Lewis equation:

$$\sigma = \frac{F \cdot p}{w \cdot Y} \tag{7.3}$$

where σ is the bending stress, F is the tangential load, p is the diametral pitch, w is the face width, and Y is the Lewis form factor.

Lewis form factors were established by Lewis (hence the name) at the end of the 19th century to account for the geometry of the gear teeth. Tables have been established for the various numbers of teeth and for the different pressure angles. For a gear with 11 teeth and a pressure angle of 20 degrees, the Lewis form factor is equal to 0.226, considering the load near the tip of the tooth. With the loads considered near the middle, Lewis factors are significantly higher (roughly 70% larger), leading to lower estimated stresses. The stress could also include Barth velocity factor to take into consideration the stresses due to impacts, once the velocity is known.

With a face width of 0.5 inch, the maximum stress on the pinion is estimated to be:

$$\sigma = \frac{750 \cdot 16}{0.5 \cdot 0.226} \approx 106,000\text{psi} \tag{7.4}$$

While this is fairly high, the stress remains acceptable. For example, quenched 8620 steel has a yield strength of 150,000psi [61]. Thus, a 16-pitch, 11-tooth, 0.5-inch-width pinion seems an excellent choice for the application considered here.

A rack with a 20° pressure angle has a Lewis form factor of 0.484 for loadings at the tip. The racks always have the highest values possible, so only the pinion needs to be considered: if it can withstand the loading, the rack will as well.

Another consequence of the pinion spinning slower is that, for a given power, more torque needs to be transmitted. The one-way bearings have to be chosen accordingly. With the gear selected and the maximum force set at 750lbf, the maximum torque is:

$$T_{\max} = F_{\max} \cdot R_{pinion} = 750 \cdot 0.34375 \approx 258\text{in.lbs} \tag{7.5}$$

The one-way bearing for 20mm (0.787in) shafts has a torque rating of 22ft.lbs, which is equal to 264in.lbs, making it the perfect candidate for this application.

The manufacturer Stieber mentions in its catalog [62] that one-way bearings can support a maximum torque equal to two times their rated value. The number of cycles that the clutch can withstand in its lifetime suffers somewhat (Figure 7.13), but 15 million cycles should be acceptable in most applications. Therefore, a one-way bearing smaller than for a 20mm shaft could eventually be used if the space available would require it.

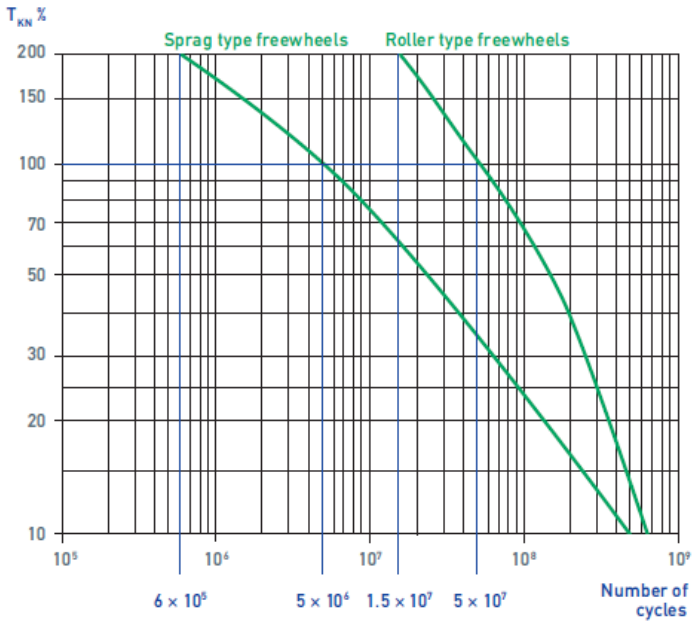

Figure 7.13: Expected lifespan of one-way bearings based on applied torque [62]

The total gear ratio from the pinions to the generators is 41.5:1. This means that the generators rotate about 19 times for every inch travelled by the rack. This is a rather high value. The goal was to determine what would be the maximum in the available space. The gear ratio can easily be reduced to the optimal value determined with the simulations. Nevertheless, "high" speeds are always desirable. This allows using generators with higher Kv values, resulting in lower internal resistances and leading to smoother outputs. But a gear ratio too high will generate high forces when the system is subjected to large accelerations. The simulation should determine the values to use for the harvester.

In order to maximize space, low friction lining to support the back of the rack was used. Nylon, Ultra-High-Molecular-Weight Polyethylene (UHMW-PE), and Polytetrafluoroethylene (PTFE), as known as Teflon (the brand name by DuPont Co.), have very low coefficients of friction: as low as 0.04 against polished steel, when lubricated [63] (Nylon is even listed with a minimum of 0.02). So, with a coefficient of friction of 0.04 and a pressure angle for the rack of 20 degrees, the friction force is:

$$\left.\begin{aligned} F_N &= F \cdot \tan 20 \\ F_f &= 0.04 \cdot F_N \end{aligned}\right\} F_f = 0.04 \cdot \tan 20 \cdot F = 0.0146 \cdot F \tag{7.6}$$

So, the choice of a friction guiding system for the rack only reduces the efficiency by less than 1.5%, which is a very good trade-off for the simplicity of the design. A rolling element could eventually be used to try to reduce friction even more, if needed. In regard to reliability, UHMW-PE is highly resistant to abrasion: it is about 10 times more resistant than carbon steel [63]. Therefore, it should not be a source of failure.

With spur gears, it is relatively easy to add a torque-limiting device. The easiest way is to have two friction discs (one on each side) pressed against the gear by Belleville washers. The gear is not rigidly connected to the shaft, but instead transmits the power through the friction discs. When the torque gets too high, the gear starts to slip. The preload on the spring washer determines the maximum torque allowable. Because of its simplicity and robustness, this system is used in many applications.

It is important to mention that care has been taken to use only elements (gears, bearings) in the design that are standard or readily available, thereby limiting custom parts to a minimum.

7.3.3 Modified Prototype

A concept using two pinions driven by a single rack is presented in Patent US 2013/0008157 [34] filed by Zuo et al. and published on January 10, 2013. The prototype presented above might therefore now fall into one of their claims: *Claim 28*:

"The energy generating device of claim 17, wherein the rack includes first rack gears;

wherein the gear assembly comprises:

a first pinion engaged with the first rack gears;

a first pinion shaft;

a first roller clutch included within the first pinion, such that the first roller clutch drives the first pinion shaft when the first pinion shaft rotates in a first direction and disengages from the first pinion shaft when the first pinion shaft rotates in a second direction opposite the first direction; a second pinion engaged with the first rack gears;

a second pinion shaft engaged with the first rack gears;

a second pinion shaft;

a second roller clutch included within the second pinion, such that the second roller clutch drives the second pinion shaft when the second pinion shaft rotates in the second direction and disengages from the first pinion shaft rotates in the first direction; and

a transfer mechanism for coupling the first pinion shaft to the second pinion shaft such that the first pinion shaft is driven in the first direction when the second pinion shaft is driven in the second direction; and

wherein the generator is driven according to the rotation of the first pinion shaft." [34]

It is worth bringing up that *Claim 17* mentioned above describes, in very generic terms, any device that uses a rack, a gear assembly rectifying the motion (which is described in more detail in the following claims), and a generator outputting direct current: "*the generator engaged with the gear assembly for receiving the rotational movement output from the gear assembly and outputting a* ***direct current*** *according to the rotational input from the gear assembly."* [34]

It would then be fair to assume that since the energy harvester presented here will be using a brushless generator outputting alternating current (AC), it is not covered by this claim.

But in order to be on the safe side and stay clear of the claims in *Claim 28*, the design has been modified. The solution is to move the rack slightly back (about one-eighth of an inch) so it no longer touches the two pinions, and to place an intermediate gear that is driven by the rack and is driving the two pinions (Figure 7.14).

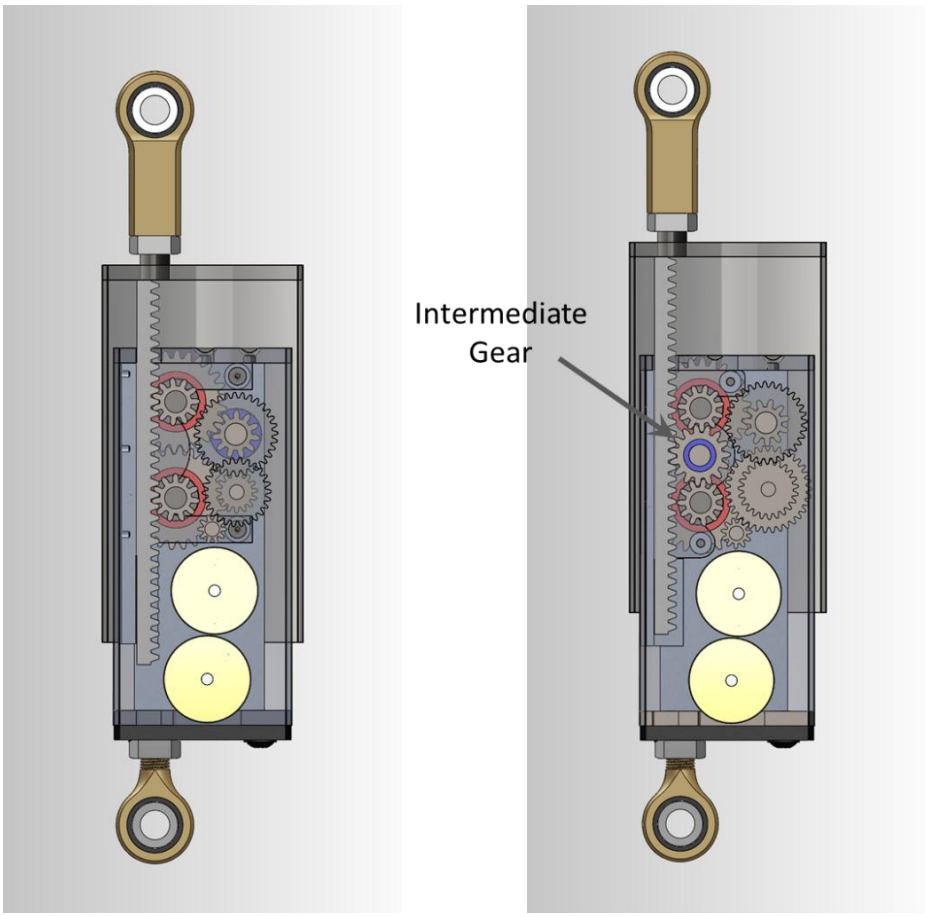

Figure 7.14: Initial (left) and modified (right) prototypes

From the two pinions up to the generators, the system remains (almost) the same. The only slight difference is that to accommodate the new gear, room has to be made and some of the gearing is changed. It can be claimed that the gearing system now relates to the application examples present in NSK and Timken's catalogs (Figure 7.10): a rack drives a pinion whose motion is the input of the illustration on the right.

There are several slight drawbacks to the modifications. There is one more gear spinning back and forth. This slightly increases the equivalent mass of the rack (+5%), which is less desirable as it increases force due to inertia. This also reduces the (maximum) gear ratio from 41.5:1 to 32.2:1, but it remains fairly high and might not be a real issue. Some of the gears have to be bigger which minimally increases the inertia: +0.2% on the equivalent inertia at the generators. There is a significant advantage resulting from the change: the possible travel increases by 0.8 inch.

The outside dimensions remain unchanged and the prototype still fits nicely inside a D5 spring (Figure 7.15). On the view from the top, it can be noted how the gears fits inside that round space.

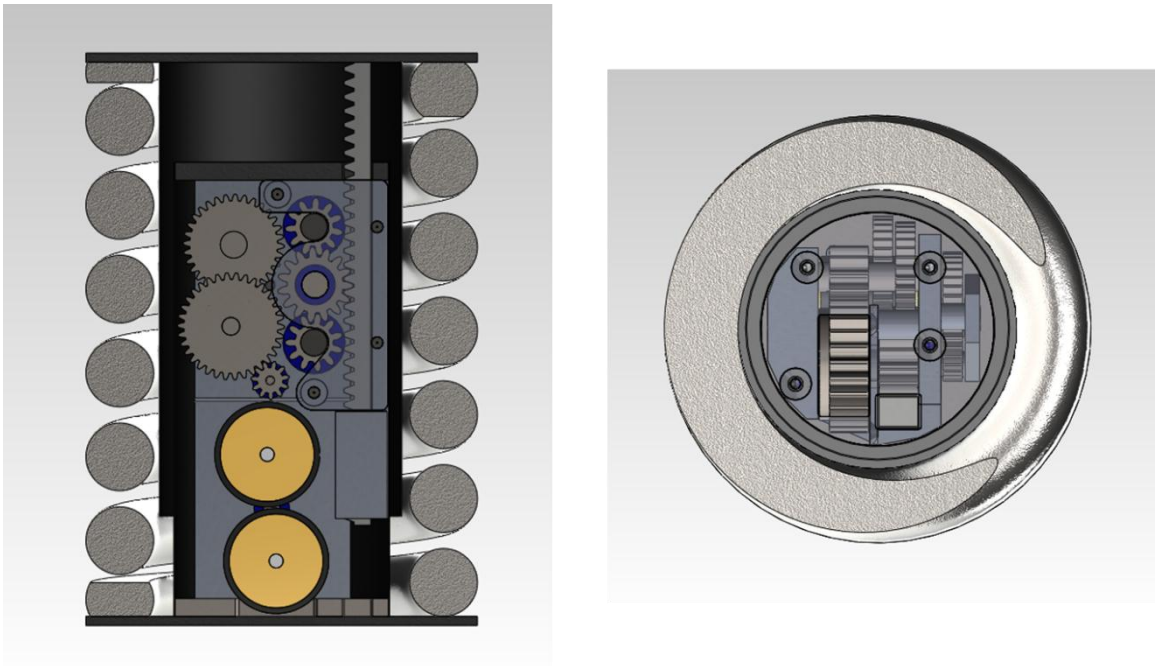

Figure 7.15: Prototype inside a D5 spring: side and top sections

7.3.4 Dynamic Modeling

The modeling of this particular concept is fairly easy. Compared to the ball screw design that has four clutch bearings and 5 kinematic groups, the new concept only uses two one-way bearings and has two kinematic groups: the rack with its pinions, and the generators with the gearing system. With two one-way bearings, there are only four combinations possible, and with one forbidden (both roller clutches engaged), there are really only three:

- if $\omega_{in} \geq \omega_{out}$, where ω_{in} and ω_{out} are the angular velocities of the pinion and the gear on the roller clutch, respectively, then one of the clutch bearings is engaged and $\omega_{gen} = k \cdot \omega_{in}$ where ω_{gen} is the angular velocity of the generator and k is the gear ratio,
- if $\omega_{in} \leq -\omega_{out}$, then the other clutch bearing is engaged and $\omega_{gen} = -k \cdot \omega_{in}$

- if $|\omega_{in}| \leq \omega_{out}$, then no clutch bearing is engaged and the generators are spinning on their inertia, slowed down by the electromechanical torque and friction: $I_{eq}\dot{\omega}_{gen} = -T_{gen} - T_f$

I_{eq} is the equivalent inertia of the generators and the gears at the generator shaft. With two gears of inertia I_1 and I_2, with N_1 and N_2 teeth, the equivalent inertia on the shaft of gear 1 is:

$$I_{eq,1} = I_1 + \left(\frac{N_1}{N_2}\right)^2 I_2 \tag{7.7}$$

This is used with the four gear stages to obtain the equivalent inertia at the generators.

Similarly, the inertias of the pinions driven by the rack can be converted to an equivalent mass:

$$m_{eq,\,pinion} = \left(\frac{1}{R_{pinion}}\right)^2 I_{pinion} \tag{7.8}$$

where R_{pinion} is the pinion radius, and I_{pinion} is its inertia (around its axis of rotation). The equivalent masses of the three gears (and shafts) moving with the rack are added to the rack mass.

All the inertias and masses are obtained with the CAD model. A Simulink model (Figure 7.16) is then built to simulate the behavior of the harvester and estimate the output voltage and the force required. The simulation is displacement-based, and the model clearly shows that the motion is transmitted through the different elements up to the estimation of the voltage, leading to a current and then a torque which moves back up the elements to give a value of the force required.

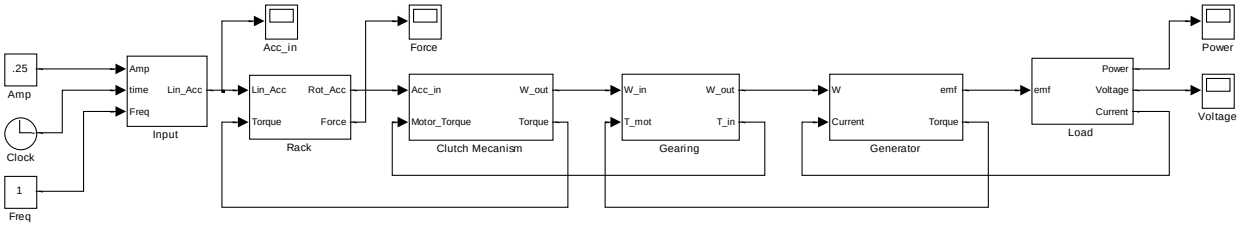

Figure 7.16: Simulink model of the rack-and-pinion harvester

With a ±0.5-inch at 2Hz input and a 50Ω load, the harvester generates 57V and 65W at peak, and at steady state, the output never goes under 30V or 18W (Figure 7.17). The effect of the gearing system with

the one-way bearing is obvious: the output is smoothed out and a battery could easily be constantly charged.

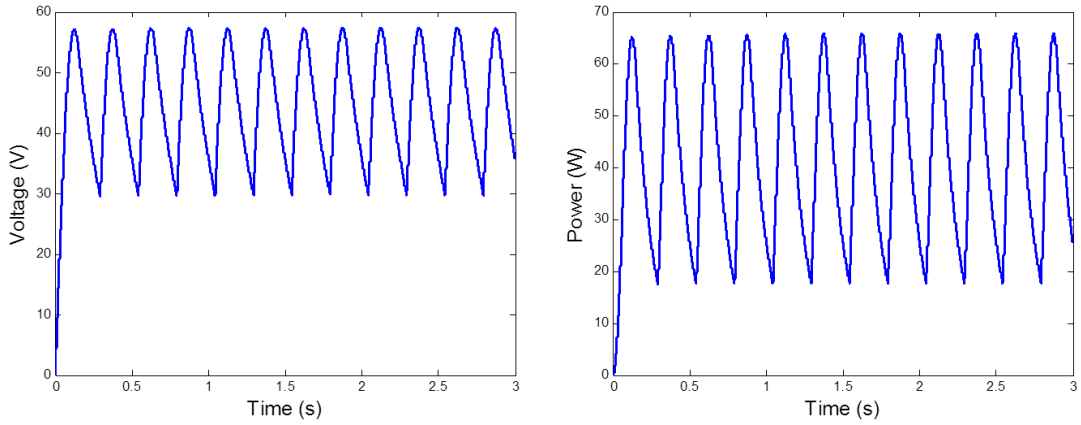

Figure 7.17: Time outputs of the rack-and-pinion harvester

The plot of the force on the rack (Figure 7.18) clearly shows when the roller clutches are engaged and that between those periods, almost no force is required (just enough to overcome friction). It can be noted that the first half-oscillation requires more force as the parts have yet to move.

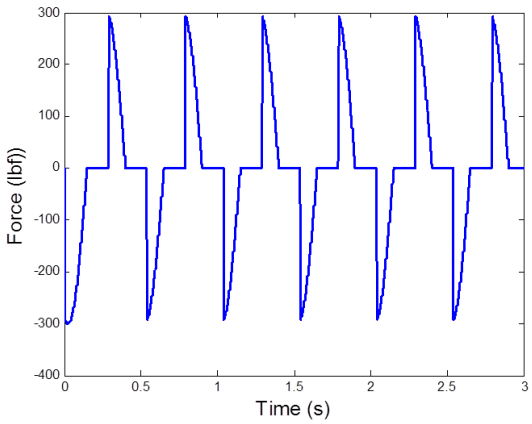

Figure 7.18: Forces on the rack to move the harvester

7.3.5 Conclusion

An energy harvester using a rack and pinion has been presented. Due to a recent patent, a small modification has been made to the design in order to be sure that it is outside the claims. With the harvester shown to be technically possible, the simulation then shows that it should also perform quite well. A prototype would be needed to validate the estimated output and performance.

7.4 Conclusion

The rack-and-pinion arrangement is the main alternative to using a ball screw, and while it comes with challenges, these can be overcome with careful and intelligent design. There are several main attractive features. One is that it is possible to have the same gear ratio in both directions. Obtaining different ratios is also achievable simply by choosing gears on the one-way bearings of different sizes. This concept only needs two one-way bearings, which simplifies the design and limits the failure modes. One last benefit to mention is that the total inertia of the parts moving in both directions with the rack is significantly lower than for the ball screw design. This can help limit the forces generated by high accelerations.

Conversely, there are several drawbacks to this design. The main problem is that only smaller generators can be used, thus not fully maximizing the space available. High stresses on the gears could be an issue, especially for fatigue. Small pinions are required, but even then, a relatively high gear ratio will also be required. The last element to consider is that the loading will be off-center. The rack cannot be in the middle, so either the mounting points are aligned with the forces (but on the side), or they are centered, where they will generate bending moments.

In summary, using a rack and pinion is an attractive alternative that would be quite interesting to prototype and test.

Chapter 8

Axle Generator

8.1 Introduction

The primary focus of the research work is to design a system to harvest wasted energy from the suspension. However, the larger picture is to find a solution to the lack of options available to power electrical equipment in rail cars. In addition, other concepts besides a damper-like device are investigated, including the axle generator. The work presented in this chapter has been presented at the 2012 Joint Rail Conference in Philadelphia, PA [64].

8.2 Design

The underlying idea of the concept is to use the rotation of the axle as input. This is not a novel concept since, as mentioned in the introduction, many systems are already using this principle. However, the design presented in this paper is unique. The system (Figure 8.1) is placed around the axle, whose motion is transmitted to the generator through friction wheels.

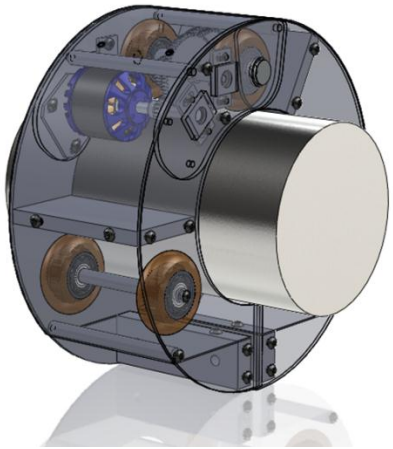

Figure 8.1: Semi-transparent 3D view of the CAD model

Three sets of two wheels are used in this design. Placed at 120° from each other (Figure 8.2), the wheels maintain the system centered around the axle. Although all three wheel sets could be driving a generator, in this first prototype only one is doing so. The other two wheels are free, with the intention of maintaining the contact between the driving wheel and the axle. From the driving wheel set, two sets of gears are used to transmit the motion and increase the speed by a factor of 8. This allows the generator to produce higher voltages, thus increasing the useful range of speed of the system. The stator of the generator is held on one of the side plates, while the rotor is placed on the last axle of the gear set.

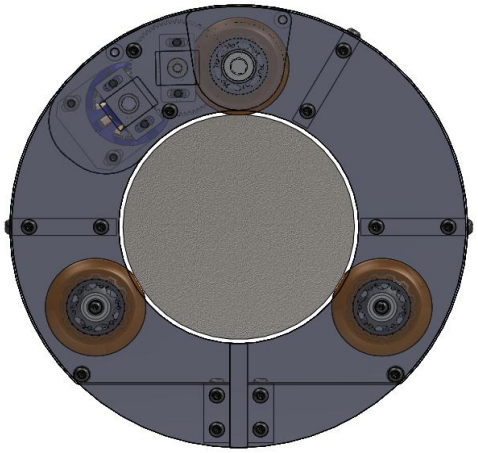

Figure 8.2: Side view (with side plates transparent)

The system is actually composed of two halves, hinged on the same axis as the axis of rotation of the driving wheels (Figure 8.3). That degree of freedom creates an opening through which the car axle can pass.

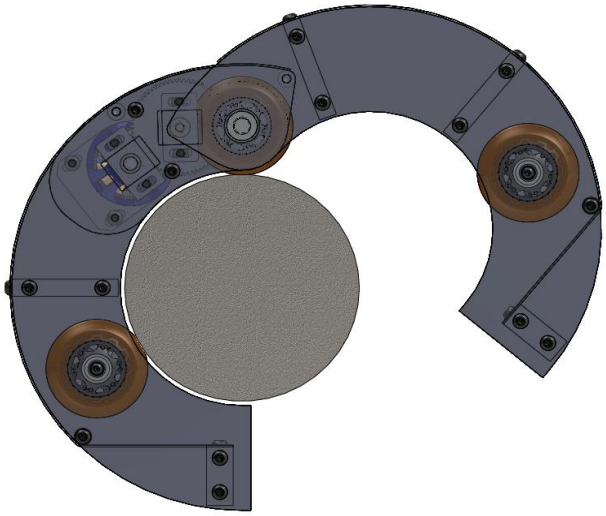

Figure 8.3: Side view - System open

So, putting the system in place is fairly easy: the halves are spread apart, the system is slid onto the axle, and the halves are closed, surrounding the axle. Spring-loaded bolts are then placed to keep the two halves together. The installation can virtually take only seconds.

The springs in the closing mechanism are there as compliant elements if a wheel or axle is not perfectly round, or if there is some debris on the rolling surfaces. They also compensate for the wear of the wheels.

To prevent the entire system from rotating with the axle, a bar needs to be secured between the system and a fixed point either on the side frame of the bogie or on the bolster. A pinned-pinned bar (hinged at both ends) would be perfect, as it prevents complete rotation but still allows some motion. It would work with the attachment point on the bolster, which moves up and down with the suspension (not shown).

8.3 Prototype

With the details of the design defined, the next step is to build a prototype (Figure 8.4).

The generator chosen for this application is the same as the one used in our suspension energy harvester. It is a repurposed permanent magnet DC brushless motor: a MonsterPower 160 from Exceed RC. Originally meant to be used in RC planes, these generators work very well as generators. Off-the-shelf components are also used for an easier and more cost-effective solution. The wheels for the prototype are those used for inline skates, and they have been chosen for their durability and grip. Their large availability was also an important factor in the choice, as it makes them easy to obtain and to replace (if necessary). Their standard size (bearing and width) and the large number of choices make the trial of different combinations (outside diameter, hardness, profile) that much easier.

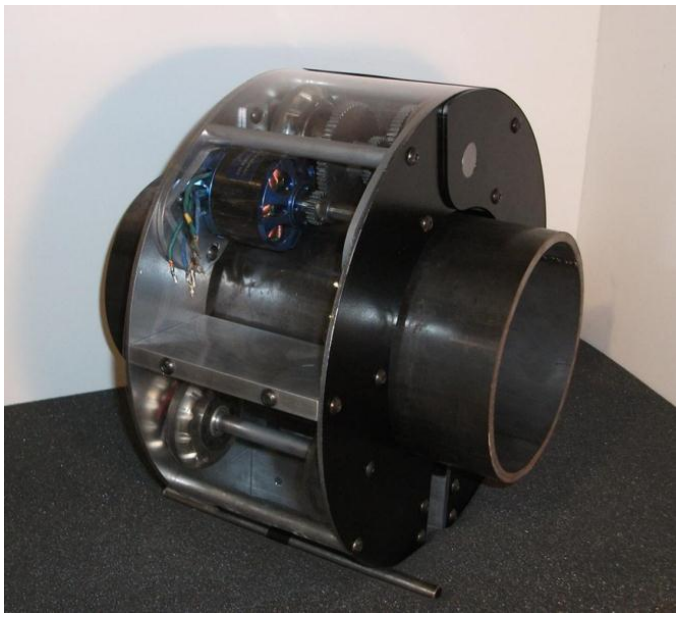

Figure 8.4: Axle generator prototype

The sides of the system have been outsourced, and are manufactured by a water-jet company. All the other parts (axles, shafts, spacers, etc.) have been machined on the lathe in-house.

8.4 Test Setup

With a fully-assembled, fully-functional prototype (Figure 8.4), initial testing and characterization of performance were completed.

All the tests thus far have been conducted in our laboratory. A 6-inch steel pipe was used to replicate the axle of the car. Close in size and material, it is a good way to approximate the real conditions of use. This simulated axle was rotated in a lathe, which is a particularly convenient way to replicate the condition of rotation of a car axle. The main limitation of the lathe is in the choice of speeds, with only eight available; the speeds range from 30rpm to 800rpm. Assuming a 36in-diameter wheel, this corresponds to 3.2mph and 85mph, respectively, which are sufficient to cover the entire range of speeds for a freight train.

The steel pipe was held on one side by the chuck and was supported on the other side by a live center. The axle generator was then placed on the pipe. A tension strap was used to prevent the system from rotating. It offers less lateral stability than a solid link, but performed well during all the testing.

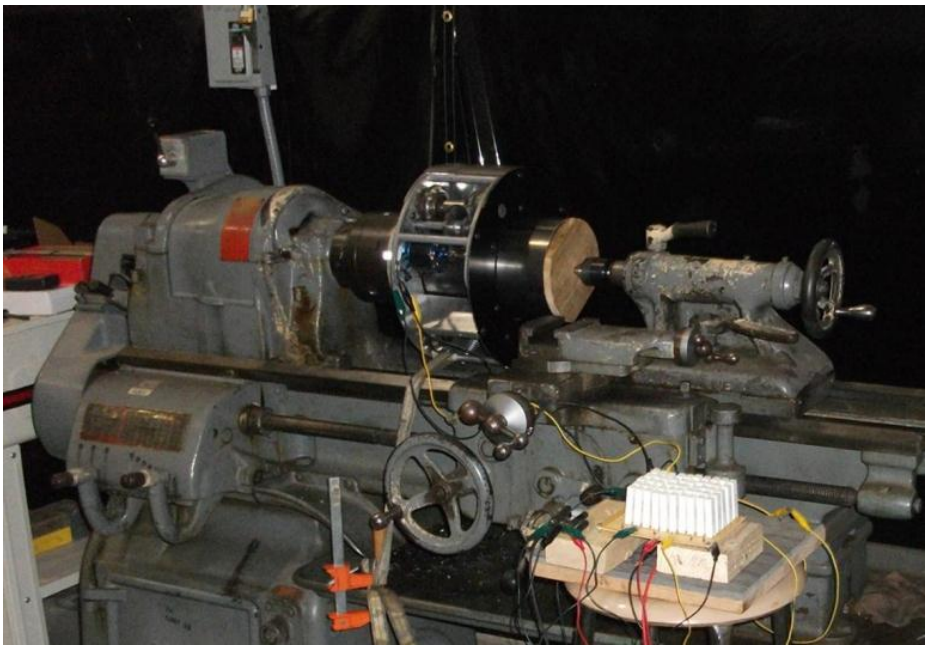

Figure 8.5: Test setup

The three phases of the generator were connected to a 6-diode bridge to rectify the voltage. The output of the bridge was connected to a power resistor, and different resistor values were used to emulate different current conditions. A voltmeter (Fluke 289 true RMS multimeter) and a digital oscilloscope (Tektronic TDS2004B) measured the voltage across the resistor. The current and power were then determined by using Ohm's law.

8.5 Test Results

For the initial tests, the prototype was tested in thirty different conditions. Six speeds were tested and five resistors were used as load. The speeds ranged from about 50rpm to 510rpm, or 5mph and 55mph, with 36-inch wheels.

The loads connected to the output ranged from 1MΩ (the internal resistance of the oscilloscope) to 10Ω. The different resistance values allowed different currents to be drawn and thus the electrical performance to be characterized under various conditions.

The Voltage vs. Speed plot (Figure 8.6) shows the linear relationship between the output voltage and the speed. We can also noted that with lower resistance, as more current is drawn, the voltages tend to be slightly lower. At 510rpm, the output voltage reaches about 35 Volts.

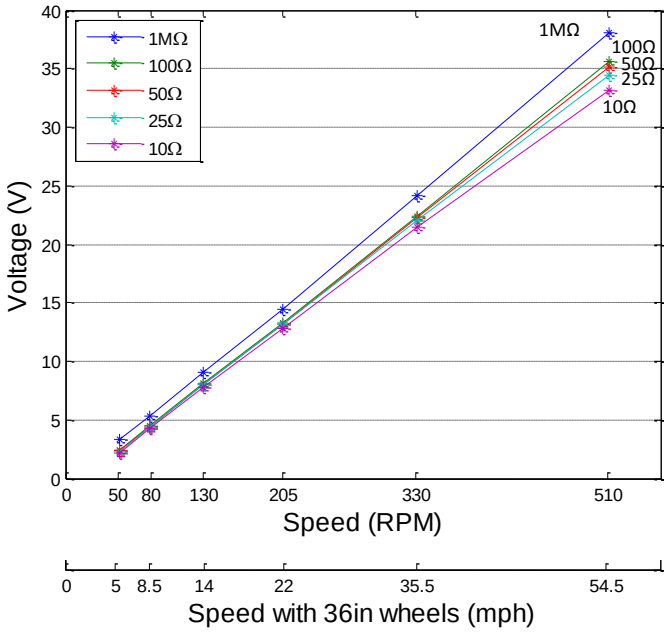

Figure 8.6: Voltage vs. Speed characteristics

With purely resistive loads, the output power is proportional to the square of the voltage:

$$P = V \cdot I = \frac{V^2}{R} = R \cdot I^2 \tag{8.1}$$

With the voltage proportional to the speed, the power is quadratic, as seen in Figure 8.7.

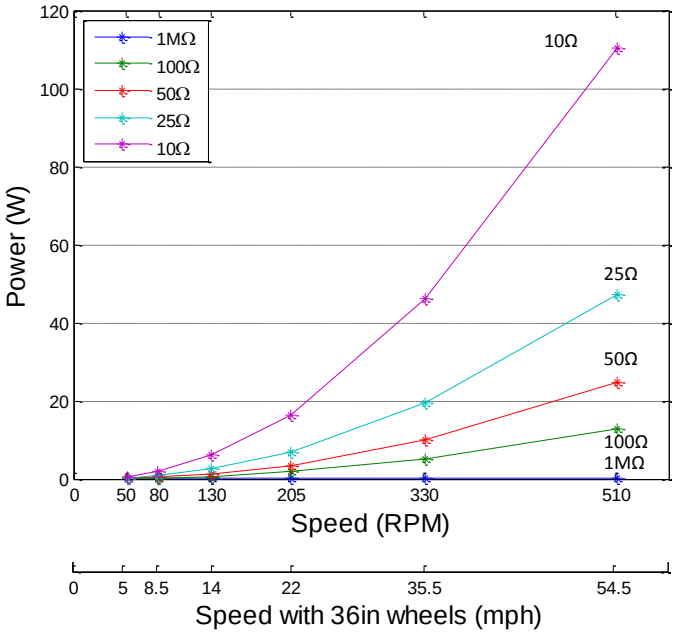

Figure 8.7: Power vs. Speed characteristics

Lower resistances yield lower voltages, but the gain in current is much higher, leading to more output power. With the resistors used during the tests, at the same speed, the result was the lower their values, the higher the power.

With the 10Ω resistor, at the highest speed (510rpm), the system is able to generate 110 Watts. That kind of power can easily supply most of the electronic systems.

8.6 Modifications and Improved Results

For the first series of tests, the generator used is the unmodified Exceed RC MonsterPower 160. From factory, its three phases are connected in a delta configuration (Figure 8.8 – left), which is more beneficial for its intended purpose as a motor. But a "Y" configuration (herein after referred to as "wye configuration"), shown in Figure 8.8, is preferable for a generator because the two phases contribute to the output voltage instead of just one. Due to a phase difference, the voltage does not double. For example, with a purely sinusoidal signal, the output is 73% higher (the exact ratio being $\sqrt{3}$). A brushless DC motor does not generate a perfectly sinusoidal output voltage when used as a generator, so the output increase might be slightly different.

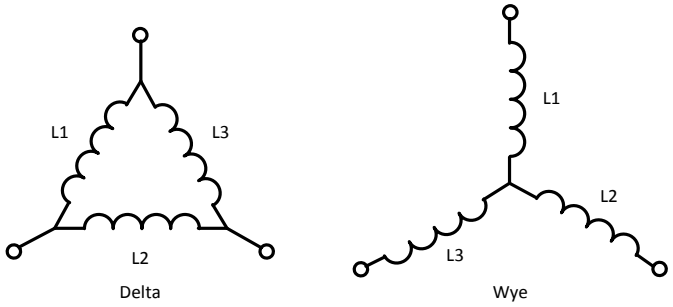

Figure 8.8: Delta and wye configurations

The connections of the phases of the generator are then changed from the initial delta configuration to the wye. The system is now expected to generate voltages about 70% higher, which would lead to three times the power.

No other modification is done to the generator (no changes in the coils themselves), and the rest of the system remains identical to that previously tested.

The system is then retested under 35 conditions; 30rpm is added to the previous 6 speeds (50, 80, 130, 205, 330, and 510 rpm). Similar resistance values are used. The tests show the same linear relationship between the speed and the output voltage (Figure 8.9).

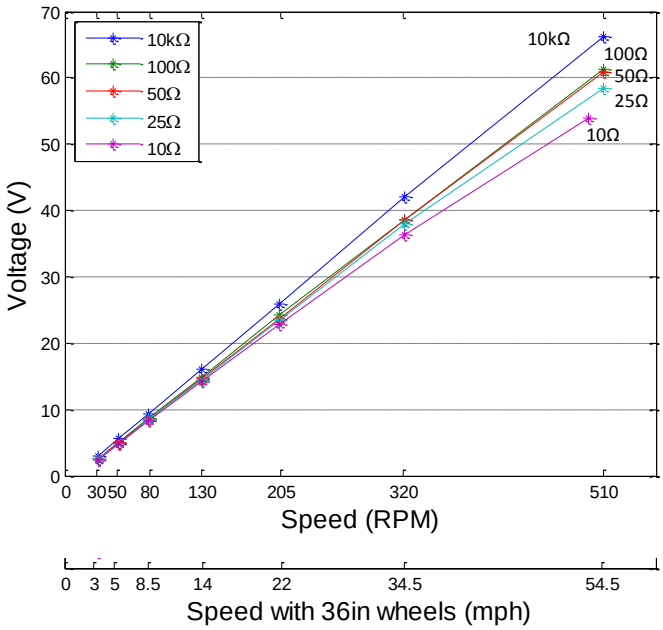

Figure 8.9: Voltage vs. Speed characteristics

The decrease of the voltages with lower resistive loads is slightly more marked, mainly due to the increased internal resistance of the wye configuration. The voltages now reach around 60 volts at 510rpm. This is a significant improvement from the previous tests. This results in an increase of about 70%, as anticipated.

The relationship between the power and the voltage is also quadratic (Figure 8.10). The recorded output powers are (approximately) three times higher than in the first series of tests. The output power is measured as high as 290W at 500rpm.

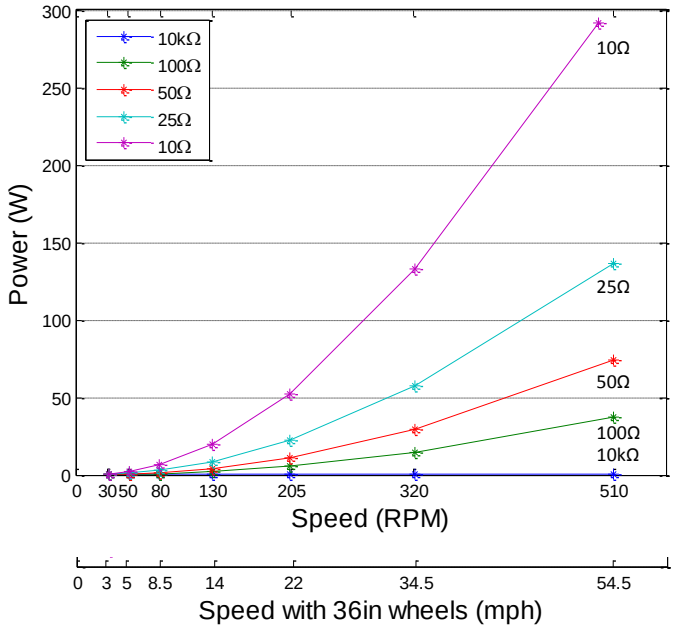

Figure 8.10: Power vs. Speed characteristics

The change from a delta connection to a wye makes a large improvement in the voltages and powers generated, allowing an even wider range of electronic devices that could be powered.

8.7 Concluding Remarks

A system for powering electrical equipment on freight cars has been presented. The design (including the concept of using the rotation of the axle through friction wheels), the development of a prototype, and the testing have been detailed.

The tests proved that the system is able to generate significant amounts of power, in particular when the generator is connected in a wye configuration. The results also show that the wheels have enough grip to transmit the necessary power to the generator. With almost 300 Watts at 55mph, the system can be used to supply power to numerous devices.

The system now needs to be evaluated for durability and to verify that the performance during extended runs remain the same as for short periods of time. Field testing is also needed to further validate the performance of the system.

This page intentionally left blank.

Chapter 9

Wayside Energy Harvester

9.1 Introduction

The main focus of the project has been on providing a solution to the lack of electrical power in freight cars, however, a similar issue is encountered on the side of railroad tracks, and the same solutions developed for onboard applications can be applied to solve this problem. Many pieces of equipment (light, monitoring systems, lubricators, etc.) are needed along the track and it is often difficult, especially in remote locations, to have a reliable source of electricity. A common solution (Figure 9.1) is to use solar panels, but besides the inherent reliance on the sun, there have been major issues with vandalism and theft.

Figure 9.1: Solar panel powered lubricator at TTC

Therefore, there is a real need from the railroads for solutions to provide power to wayside equipment. The electromechanical devices presented in this chapter are designed to respond to that need. They use, as input motion, the vertical deflection of the track due to the weight of passing cars. Since that linear motion is similar to that of a car suspension, energy harvesting systems can be developed based on the same concepts presented in previous chapters.

9.2 Existing Systems

Some solutions have been proposed by different companies and universities. Innowattech [65], for example, is an Israeli company that developed large-size piezoelectric elements that can be embedded in the road (Figure 9.2) or in the railroad ties to generate electricity when wheels pass over. They have also developed a piezoelectric pad that instead of being embedded in a custom-made tie, can be placed between a conventional tie and the rail in place of the rubber pad that is often used to dampen vibrations. Other concepts to harvest energy from railways are presented in a patent filed by the company [66]. Capability for significant power generation is claimed on their website, but no technical data has been found to verify those claims.

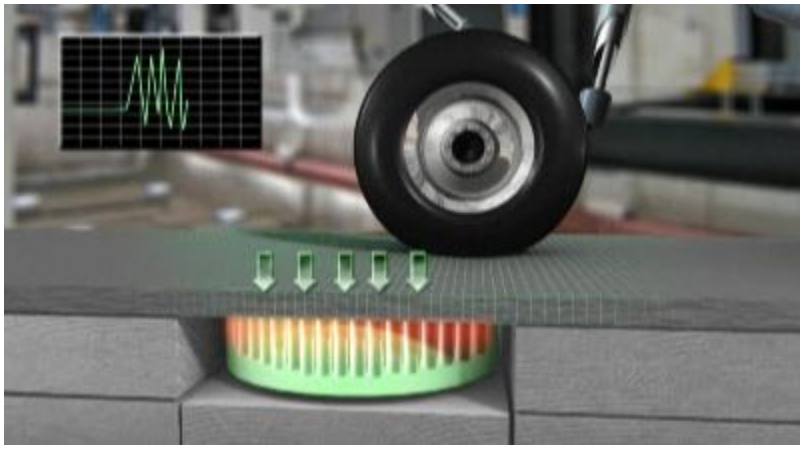

Figure 9.2: Innowattech's concept of piezoelectric energy harvesting system embedded in the road

Dr. Nelson and his team at Nebraska-Lincoln have also investigated using piezoelectric material to harvest energy from passing trains [67]. The idea is to convert the strain at the bottom of the rail (Figure 9.3) into electrical energy. Only very limited amounts of power (milliWatts) are generated.

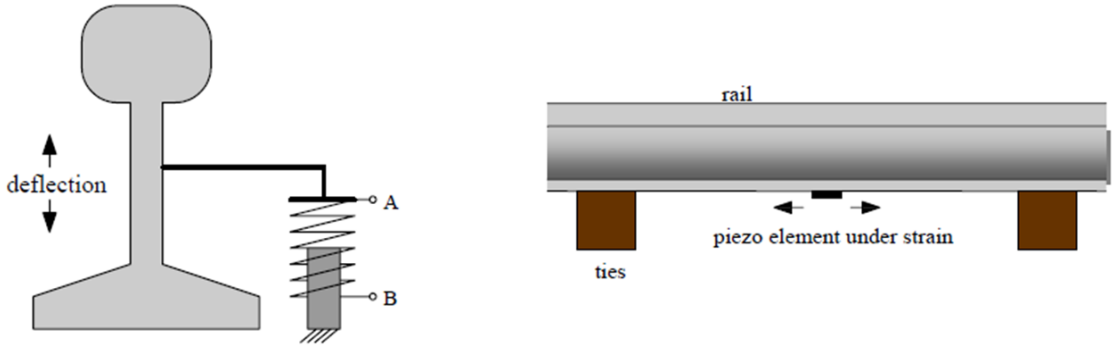

Figure 9.3: Energy harvesting systems investigated by a laboratory at Nebraska-Lincoln

They also developed a prototype using magnetic induction. A conductive coil is attached to the rail and a magnet, inside the coil, is fixed to the ground. As the track deflects under the weight of the cars, the magnet moves inside the coil, generating a variable magnetic field that in turn induces a voltage across the coil. During testing, less than a volt is generated, yielding only a fraction of a milliWatt. With heavier rail cars, it is expected to generate about 1mW [67]. To reach higher power levels, their focus shifted to a system using a rack-and-pinion arrangement to transform the linear vertical motion of the rail into rotation. A DC generator is then connected and driven by a gearbox used to increase the angular velocity of the pinion. Lab tests estimate that up to about 4Watts could be generated by this system [57], but during field tests, average powers did not exceed 0.22W [68]. Up to that point, the system only harvests energy during the downward motion. A clutch bearing allows the generator to keep spinning as the track starts to move back up. The system is then modified to harvest energy from both the downward and upward motions. Two clutch bearings are now used to rectify the motion and allow the generator to always spin in the same direction (Figure 9.4).

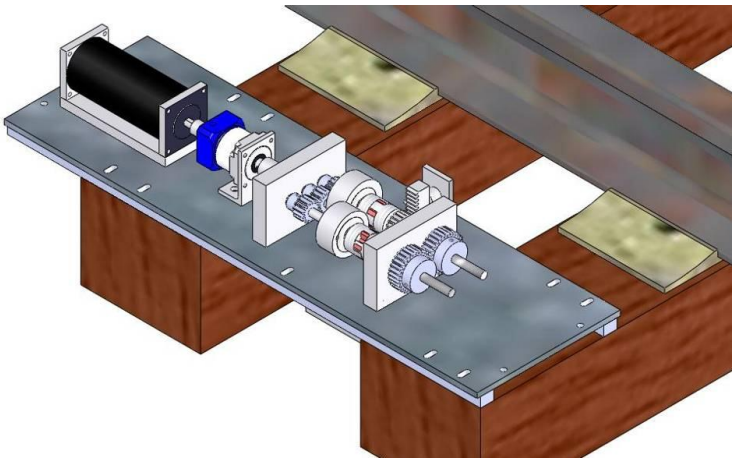

Figure 9.4: Upgraded energy harvester from Dr. Nelson et al. from Nebraska-Lincoln [58]

Recently, Dr. Zuo and his team at Stony Brook University presented their work on a similar concept (Figure 9.5). Racks are fixed to the ground and the system is attached to the ties. When the track is deflecting under load, the racks move with respect to the system, creating a rotation that is transmitted through a gearbox to the generator. Like the previous system, the motion is "rectified": upward and downward motions generate a rotation of the generator in the same direction. The preliminary prototype showed during lab tests that it can generate up to 1.4 Watts [59], but the second prototype is projected to be able to generate 200 Watts.

Figure 9.5: Energy harvester from Stony Brook University (Dr. Zuo et al.) [69]

9.3 Proposed Concepts

Instead of fixing the energy harvester across ties as others suggested, the concept is to completely replace one tie by a custom tie housing an energy harvesting system. This allows us to fully utilize the weight of the car instead of putting a lot of effort into compressing the ballast. To keep operations safe and remain transparent to the other components, the deflection of the track with the special tie needs to be controlled to remain similar to that at any other point on the track.

Different concepts are then developed, from direct implementation of the suspension harvester to designs tailored to the tie constraints.

9.3.1 Direct Implementation of the Suspension Energy Harvester

With promising results in 2010 for the new prototype for freight car suspension, many ideas for implementations where linear motion is present flourished. One of them is to place the harvesting device, unmodified, directly under the rail (Figure 9.6). It rests on a solid foundation, which could be a concrete slab, fixed with the ground. On the other end, the ball screw is attached to the rail. The tie is made out of a compliant element that can compress under load, for example, alternating layers of rubber mats and steel plates. Its composition is important to ensure a controlled deflection and give the rail the required support. With the weight of the car over the system, the tie compresses and the relative motion between the rail and the foundation is the input to the harvester: the ball screw is pushed down, and generates a rotation that is increased before being transmitted to the generator. The upward motion is also used to drive the generator. With the gearing arrangement, the generator is, of course, as explained for the suspension application, always rotating in the same direction to improve efficiency.

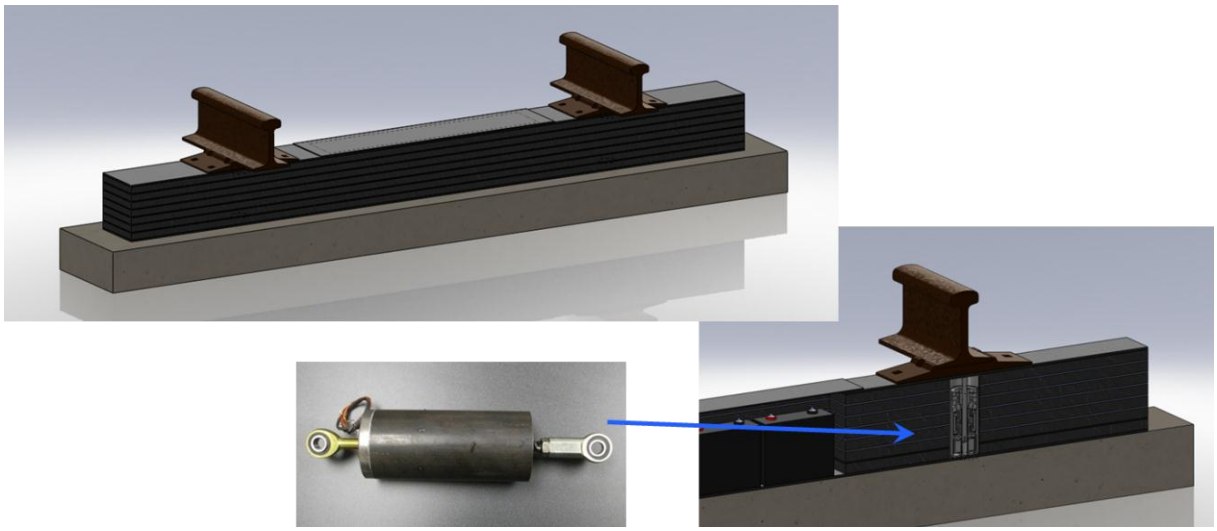

Figure 9.6: Suspension harvester to harvest energy from vertical track motion

This concept can come in different variations (Figure 9.7). A coil spring can be placed in parallel to the harvester in order to give more support and more control over the compression of the tie. Furthermore, with the room available, two harvesting systems can be placed under the rail, to take better advantage of the forces available and thus increase the power that can be generated.

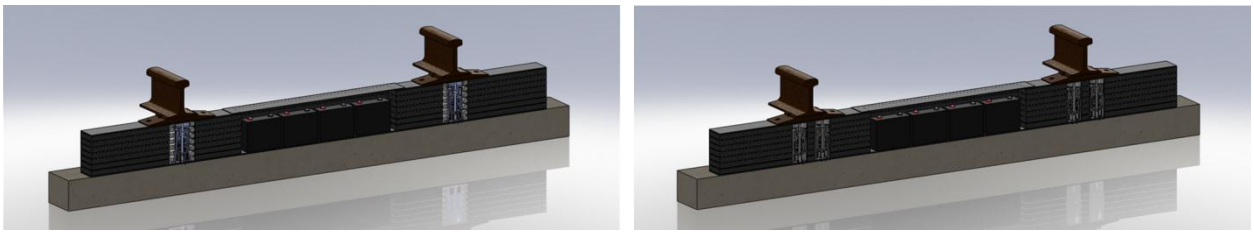

Figure 9.7: Variations of the concept using the suspension harvester without modifications

As can be seen on the illustrations, the center of the tie can be used to house some batteries and the electronics. Keeping the elements hidden inside the tie helps to prevent vandalism and theft that are an issue with this type of equipment.

Alternatively, the harvester could be flipped, with the ball screw pointing down, and fixed to an anchor point planted into the ground (Figure 9.8). The main body of the harvester is rigidly connected to the rail. A layer of compressible material is placed at the bottom of the tie to give the appropriate deflection.

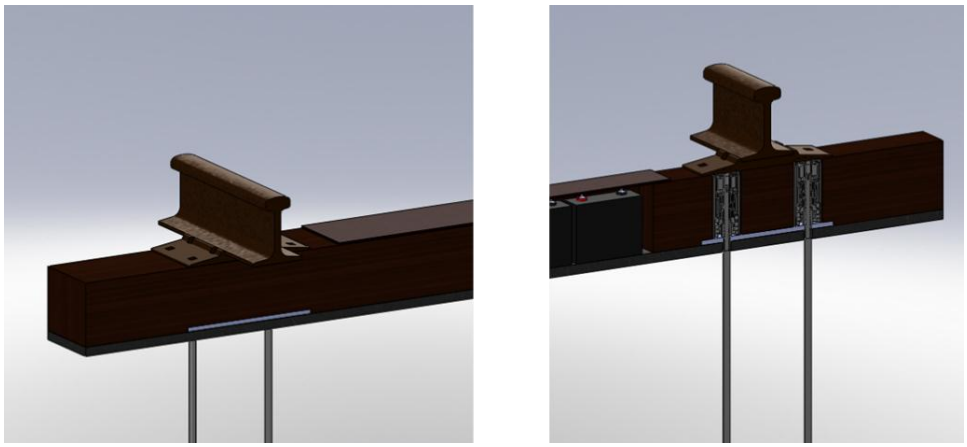

Figure 9.8: Tie with harvesters upside down (ball screws pointing down)

9.3.2 Crosstie Ball-Screw Harvester

Realizing that the tie does not have the same size constraints as a freight car suspension, the harvester can be made significantly larger. The concept remains similar, with a ball screw transforming the vertical deflection of the track into rotation, and one-way bearings creating a unidirectional motion for the generator. However, every element can be made larger (Figure 9.9) using the length of the tie. The rotating parts are no longer all concentric. A consequence is that it is now easier to obtain the unidirectional motion of the generator. Two spur gears are driven by the ball nut through two one-way

bearings such that one transmits torque during the downward motion of the screw, while the other transmits torque during the opposite motion. With a different number of external contacts, they both drive the generator in the same direction.

Figure 9.9: Crosstie ball-screw harvester

In this design, the ball screw is fixed to the ground and the rest of the harvester is moving with the rail. Conversely, the ball screw could move with the rail and the generator/gearing system could stay fix.

9.3.3 Custom-tie Harvesting System

The previous concept takes better advantage of the size of the tie than just implementing the suspension harvester. But the idea can be taken further by having the generator rotating around an axis parallel to the great length of the tie. This way, the generator is not as limited in size. With a rotation around that direction, a design using a ball screw would require some bevel gears. However, other systems are capable of directly generating the rotation in the desired direction, and, of course, the first that comes to mind is a rack and pinion system. Thus, a new concept is developed around this system (Figure 9.10). The rack is anchored to the ground and the rest of the system is moving with the rail. Two one-way bearings are again used to obtain the unidirectional motion of the generator. A flywheel is added to the generator to smooth out the angular speed and thus the output. It is easy to see that this layout allows a larger generator to be used.

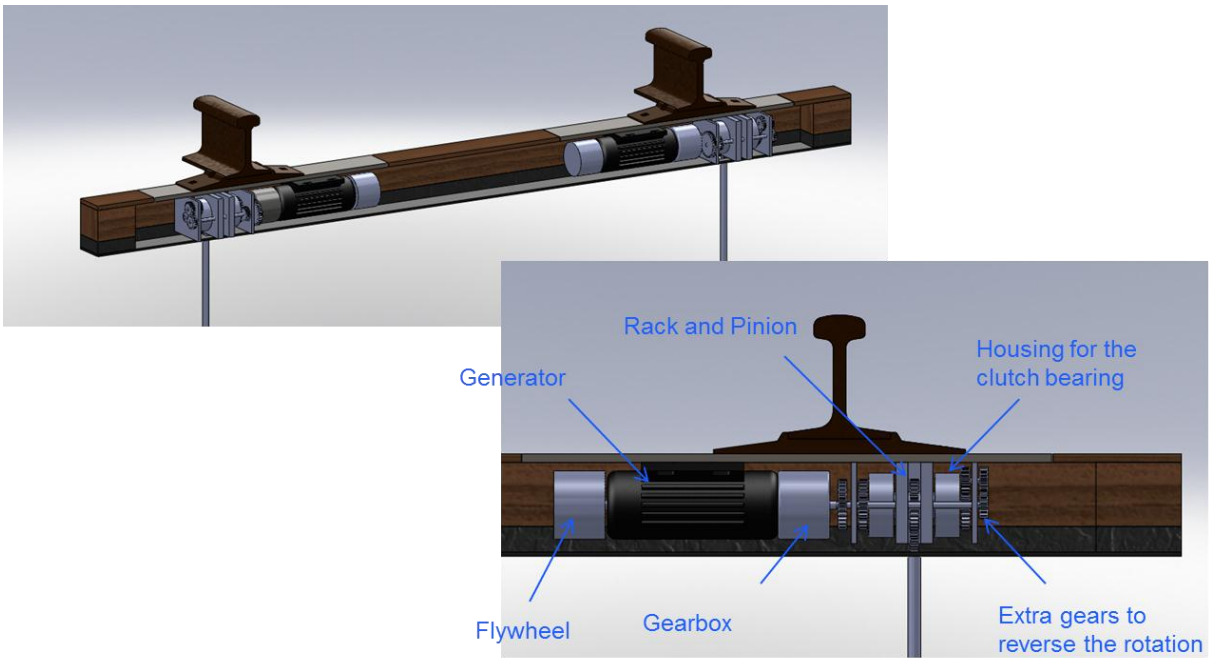

Figure 9.10: Crosstie rack-and-pinion harvester

It can be noted that the rack is slightly offset and not directly under the head of the rail. This is meant to facilitate installation, in that the rack can be reached simply by removing the baseplate.

Without considering the aspect of being incorporated inside the tie, this concept is similar to what the teams from Nebraska-Lincoln and Stony Brook University have since presented at different venues.

Besides using a rack and pinion system, solutions with a scotch yoke, a crank mechanism, or chains/straps have also been investigated.

9.4 Testing

Validation and estimation of the performance of the prototype would be extremely important, so it is essential to have a good method of testing. The Railway Technologies Laboratory (RTL) has been donated a 40-foot section of track (Figure 9.11), which could be an excellent instrument to test an energy harvester using track deflection.

Figure 9.11: 40-foot track panel at RTL

One approach is to support the harvesting tie and some of the neighbor ties by ballast, such as they would be on a regular track. To reduce the amount of ballast needed, part of the track can be supported at the correct height by solid elements like cinder blocks and lumber (Figure 9.12). If they are far enough from the harvester, they should not impact the track deflection. Two hydraulic actuators apply the normal load to the rails through an I-beam. Steel beams also connect them under the ballast, effectively compressing the track. Small compression elements are placed between the I-beam and the rails to apply the load exactly where desired. This setup allows us to conduct force-driven scenarios. The track deflection does not need to be assumed, like it would be with displacement-based tests, which contributes to making the tests more realistic.

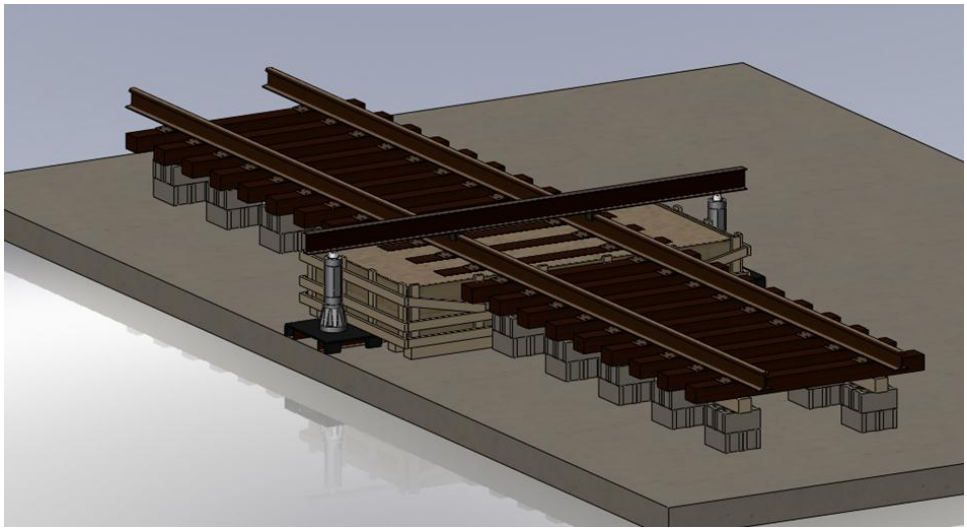

Figure 9.12: Test setup using the 40-ft track panel

Using such a long piece of track might be relatively difficult, so a more practical solution is to use only a short section (Figure 9.13). The idea is relatively close: the same loading system with two hydraulic actuators replicates the weight of the passing cars, but all parts of the track that is not supported by the ballast are now removed.

Figure 9.13: Test setup using a short track section

This setup can, like the first one, perform force-based tests. A simpler test rig (Figure 9.14) would consist of mounting the harvesting tie on top of the hydraulic actuators and fixing the rack or ball screw on the cylinder shaft. The tie would be held stationary, and only the rack/ball screw would be moving, probably hidden in the mount. This setup could be used for force-controlled tests, but since the impact of the rail and the interaction of the adjacent ties are not present, it only presents limited interest. Thus, this rig is more oriented toward displacement-based tests. They are less realistic, as the track deflection has to be assumed, but with a good model or with measurement data (assuming the system would let the track react the same way), it can replicate quite well what will be seen in the field and therefore give excellent insight into the performance of the system.

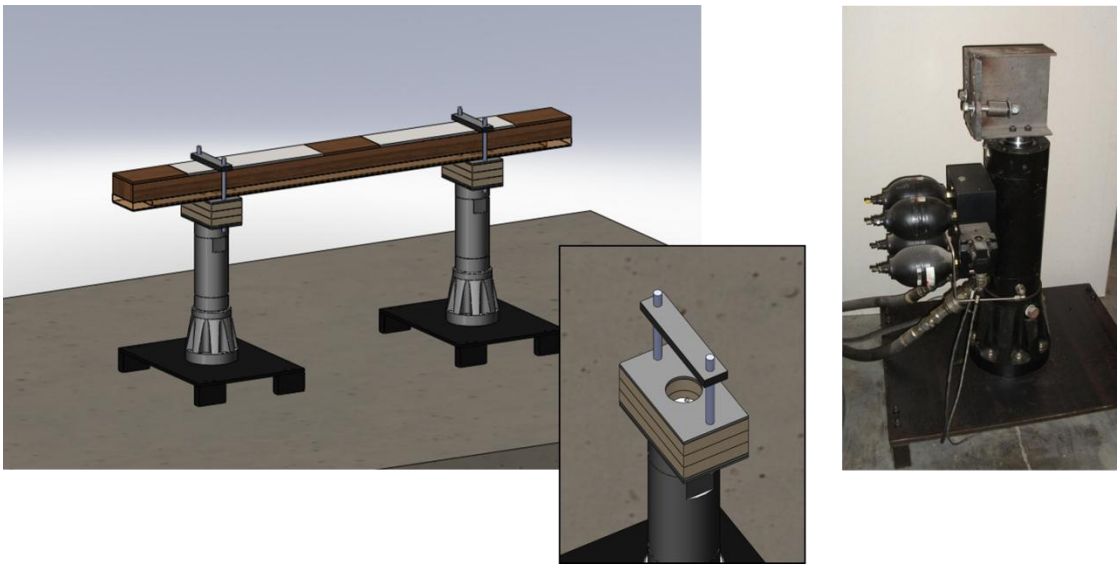

Figure 9.14: Setup for displacement-based tests

9.5 Conclusion

This chapter presented various concepts to harvest energy from passing trains to power wayside equipment ranging from the direct implementation of the system developed for freight car suspensions, to custom-designed systems that fit nicely inside a tie. The latter certainly offers more potential, as it takes full advantage of the space available instead of being limited by the smallest dimension of the tie - its height. Railroads have a real need for such a system, and for alternatives to solar panels that are often a management nightmare between vandalism and thefts. Other groups have also identified that market and have already developed solutions and built prototypes, sometimes with mixed success. So far no one has proposed to entirely replace one crosstie. Certainly, care has to be taken to maintain the integrity of the track and keep the vertical deflection within reasonable limits, but these design constraints are not insurmountable, and this system should lead to better performance with higher harvested energy.

This page intentionally left blank.

Chapter 10

Conclusion and Future Work

This chapter presents various thoughts and conclusions about the research work achieved and different directions that could be investigated in the future.

10.1 Concluding Remarks

Various energy harvesting devices were designed, prototyped, tested, and developed. The first system introduced uses a linear generator. A magnet arrangement moves linearly inside a coil, generating a voltage without requiring any mechanical transformation. Although the principle of increasing efficiency by avoiding mechanical transformation is valid, low-speed conditions are not optimal to generate electricity. Having a large coil allowed the generation of significant voltages even with small motions but the internal resistance and the losses prevented large output power. Smaller coils yielded more electrical current, but the voltage was limited by the size of the winding and the number of coils in series, still leading to limited output power. It would be possible, however, to increase the generated power by including more coils.

Converting the linear motion of the suspension into a rotation allowed the use of a gearbox that increases the angular speed. It was then possible to use a more efficient generator in which, unlike a linear generator, all the magnets contribute to generating a voltage. The core can be laminated to limit eddy currents. A simple hand-pulled prototype using a rack and pinion proved that it can generate more power than a linear generator, leading the way to a second generation of energy harvester.

Packaged like a damper, the new prototype uses a ball screw to obtain a rotation from the initial linear displacement, and the generator is connected to a gearbox directly driven by the ball nut. The tests showed that the output voltages were only marginally higher than with the linear generator, but the power was incomparably larger with up to 50 times more, as a lot more current can be drawn. A special mechanism using four one-way bearings and the first stage of the planetary gearbox led to higher efficiencies, as the generator was then only rotating in one direction and was able to rotate freely as the input slows down. With this modification, the generated power improved by approximately 50%. A key element of the harvester is obviously the generator, so it is important that it be well-adapted to the application to maximize the performance. The windings were redone to increase the output voltage while

at the same time removing one stage of the gearbox. Although spinning five times slower, the generator then managed to produce about 15% higher voltage and about 25% more power on average.

One goal was to manage to fit the system inside the spring of the suspension. A new prototype, based on the same concept but in a shorter package, was then designed. It outperformed the first prototype because of its larger size, generating more than 12 V_{RMS} with only a 0.25-inch amplitude at 1Hz. At higher velocities, it is capable of producing average powers in excess of 80 Watts. Reliability is also a key aspect of the proposed design. The prototype was tested for two million cycles, with a sinusoidal input of 0.5-inch amplitude at a frequency of 1Hz. The harvested power was dissipated through a 50Ω or 20Ω resistor. Except for minor issues, no failure that required replacing the harvester was experienced.

The harvester was also tested with suspension relative velocities that had been measured in the field. The tests indicated that the harvester can generate a few watts of power on average, and over 50 Watts at peak at usable voltages and currents. The successful operation of the harvester in conditions that replicate the field environment significantly raised our confidence in the suitability of the system for actual field testing. Additionally, it allowed us to adjust the resistive load to further improve the harvested power.

One important observation made during the tests replicating actual data was that there are high accelerations that lead to fairly high forces. This means that a safety mechanism limiting the torque is required. A solution was designed and its viability was proven. Using spring-loaded elements, the ball nut can automatically disengage when the torque surpasses a threshold, thus limiting the maximum torque transmissible. Reliability was an issue with the first prototype of torque limiter, but using a harder material and bearing balls solves this problem.

Although laboratory tests were promising, the energy harvester still needed to be tested in actual field conditions to truly validate its performance. Several prototypes were then built, integrating the various modifications brought about by the evaluation tests. Working with the Transportation Technology Center, Inc., two prototypes have been mounted on one of their test cars, and some tests have been conducted on the High Tonnage Loop. The limited test data that we have has proven that the harvesters can withstand the harsh conditions of a freight car suspension and can generate power, in the range of, and sometimes higher than, that estimated with the lab tests.

The development of a numerical model to closely match the test data and give a trustworthy estimation of the output voltage and power for any input condition was also completed. The model can be used in two ways: 1) as a tool for optimization and parametric studies, and 2) as an element in a complete suspension model to estimate its impact on the whole dynamics of the railcar.

An alternative solution to provide onboard power was presented, using the rotation of the wheelset instead of the suspension motion. In this design, a number of friction wheels roll on the axle to drive a generator. It proved in laboratory tests that it can produce large amounts of power, on the order of 300 Watts at 500rpm, which corresponds to a speed of about 55mph.

This research led to the development of feasible systems that can provide power to railcar-mounted electrical components in a reliable and cost-effective manner.

10.2 Future Work

The concept of harvesting energy from a freight car suspension with a system using a ball screw has come quite far, up to the point of being currently tested on an actual rail car. This marks a major milestone but it is not the end of the journey. The system works, but it can be improved to generate more power. With the data collected on the freight car, it is possible to correlate the lab tests in order to obtain a precise idea of the potential of the system and ultimately to tune it for maximum performance. Some aspects are directly manageable, like choosing the appropriate Kv value and the corresponding winding configuration or deciding what capacitance and equivalent resistance combination generate the most power. Another aspect to consider is using a different lead for the ball screw. A larger lead would reduce the angular speed for a given displacement, but it would also reduce the equivalent mass of the screw and thus could decrease the peak forces due to inertia. This is tuning the existing prototypes for highest performance; some other tweaking would be more of a redesign. The ball nut could be connected through its sides instead of using the thread. This would allow for the gain of some needed space inside the harvester. Also, having the ball screw in the center of the harvester requires elements of larger size which in turn are rated for torques higher than needed. The two main examples are the one-way bearings and the planetary gear that are rated for 49ft-lbs and 85ft-lbs, respectively, whereas the force allowable by the ball screw leads to a maximum torque of only 6.6ft-lbs. Most likely, it is not possible to have custom-made bearings or really flat gears, but such a configuration can optimize the design, reduce the inertia of the rotating parts and leave more room for a larger generator. The space that would be freed could also be used to include a spring element that would store some of the energy during sharp peaks and then release it progressively, smoothing the output voltage.

To move the project even further, it is becoming necessary to start targeting specific applications and collaborating with research teams and companies that want to deploy electronic devices and are looking for a reliable source of electricity. A promising possibility for collaboration would be with the group from

the Railway Technologies Laboratory, here at Virginia Tech, working on the development of an onboard rail breakage detection system, with the possibility of the required sensors being embedded in the harvester itself. Associating the energy harvesting systems with a particular system could accelerate the diffusion of the system.

At the beginning of the design of the system using a rotating generator, the choice was made to use a ball screw to convert the linear motion of the suspension into rotation. The other option would have been to have a rack and pinion system. Although such a mechanism would have some challenges in terms of generating a rotation that is perpendicular to the axis of the system, it has some advantages compared to a ball screw. It makes the rectification of the motion a lot easier to achieve, and implementing a torque limiter in spur gears is relatively simple. Therefore, the possibility of using a rack and pinion in a suspension energy harvesting system has been investigated. A complete design has been realized, and the remaining step is to build a prototype. Simulations show that it could generate a useful amount of power, and tests are now needed to confirm it.

While the railroads are seeking a solution to the lack of electrical power in their freight cars, they have a similar issue with trackside applications. Sometimes, in remote locations, there is the need for signals, measurement equipment, or rail lubricators in out-of-the-way places with no convenient power supply. The most common solution is to use solar panels, but they have different disadvantages. Besides the obvious constraint that they can only produce power during the day, the main issue reported by the railroads is that they are often stolen or vandalized. The ideas to harvest energy from the linear motion of the suspension are applicable to harvesting energy from the vertical deflection of the rail when the train passes as well. Although the track is fairly rigid, the large weight of a freight train can create up to an inch of rail deflection. Coupled with the large forces involved, there is potential to generate a large amount of power, enough to solve the issue of supplying electrical power in remote locations. Different concepts have been presented. Choosing one and creating a detailed design would be the next step, followed by the realization of a prototype and testing.

Another direction that can be followed is to investigate the possible implementations of small-size harvesters using ball screws. The design of a harvester that can fit inside a leg of a mountain bike fork is an example. This concept could be developed into a fully functional prototype and integrated in a showcase *Hokie bike*. This particular application serves as a demonstration that the freight car harvester concept can be scaled down; many other applications could be envisioned.

References

[1] Sheng Lu, "Real-time vertical rack deflection measurement system," University of Nebraska-Lincoln, PhD Dissertation, 2008.

[2] Ram Samy, Orestes Varonis, and Michael French, "Antifriction Bearing Capable of Generating Electrical Energy," 5,440,184, April 12, 1994.

[3] Mark Edwards, "Improving safety with on-board monitoring and control," in *International Conference on Communications*, Seoul, South Korea, 2005.

[4] Timken, "U.S. Federal Railroad Administration Selects Timken 'Smart' Technology," *Rail Bearing Journal*, no. 13, p. 2, 2004.

[5] Michael French, Khan Iftekharpuddin, and David Leeper, "Bearing with Wireless Self-Powered Sensor Unit," 6,535,135, March 18, 2003.

[6] "Bearing Down," *Progressive Railroading*, Sept 2006.

[7] Tieme Stikkers, "Drive Apparatus for Vehicle Mounted Generator," 4,165,466, Aug 21, 1979.

[8] John Scott and Dudley Peters, "Axle Mounted Alternator for Railroad Cars," 4,377,975, Mar 29, 1983.

[9] Axel Kemmer, Mahmud Keschwari, and Peter Zeller, "Generator Arrangement for Producing Electrical Energy and for Use in Vehicle with Anti-skid Systems," 5,585,711, Dec 17, 1996.

[10] John Barrett, "Railroad Car with Axle-Mounted Electrical Generator," 5,828,135, Oct 27, 1998.

[11] Phil Langhorst, Donald Merz, and Paul O'Connor, "Generator Mounting Assembly," 5,915,306, Jun 29, 1999.

[12] Kenneth Folk and Donald Good, "Apparatus and Methods for Mounting a Generator to the Axle of a Railcar," 6220175, Apr 24, 2001.

[13] Wayne Murray, "Self-Regulated Axle-Mounted Electrical Generation Device," 6,474,832, Nov 5, 2002.

[14] Robert Martina, Michael Asher, James Brannon, and Carl Chance, "Alternator System for Rail Car Electronics," 2006/0169169, Aug 3, 2006.

[15] James Wayland, "Generating Electricity for Cars," 627,246, June 20, 1899.

[16] James McElroy, "Device for Attaching Dynamos to Railway Cars," 731,382, June 16, 1903.

[17] Winston Brecht, "Axle Generator Drive," 2,032,112, Feb 25, 1936.

[18] William Whitsitt, "Drive Gearing for Axle Generator," 2,053,983, Sept 8, 1936.

[19] James Last and David Rowe, "Method of Generating Electricity and Electrical Generator," 3,696,251, Oct 3, 1972.

[20] John Konotchick, "Linear Motion Electric Power Generator," 5,347,186, Sept 13, 1992.

[21] Pat Mah, "Faraday Flashlight," 6,729,744, May 4, 2004.

[22] http://www.milesaheadnetwork.com.

[23] Oskar Danielsson, "Design of a Linear Generator for Wave Energy Plant," Uppsala University School of Engineering, Uppsala, Sweden, Master's thesis, ISSN 1401-5757, 2003.

[24] Fred Gardned and Rudolfus Van Schie, "Device for Energy Conversion from Wave Movement," 6,256,985, July 10, 2001.

[25] M. Leijon et al., "An electrical approach to wave energy conversion," *Renewable Energy*, vol. 31, p. 1309, July 2006.

[26] Lorand Szabo, "Novel permanent magnet tubular linear generator for wave energy converters," in *2007 IEEE International Electric Machines and Drives Conference*, Antalya, Turkey, 2007, pp. 983-7.

[27] Irina Ivanova, O. Agren, H. Bernhoff, and M. Leijon, "Simulation of a 100 kW Permanent Magnet Octagonal Linear Generator for Ocean Wave Conversion," Uppsala University, Uppsala, Sweden,.

[28] Henk Polinder, Michiel Damen, and Fred Gardner, "Design, modelling and test results of the AWS PM linear generator," *European Transactions on Electrical Power*, vol. 15, p. 245, May 2005.

[29] http://www.perpetuum.com.

[30] Shakeel Advadhany, Paul Abel, Vladimir Tarasov, and Zack Anderson, "Regenerative Shock Absorber," 2009/0260935, October 2009.

[31] Levant Power. [Online]. http://www.levantpower.com

[32] Shigeru Yamawaki, Yasuo Shimizu, Katsuji Watanabe, and Atsuhiko Yoneda, "Electric damper," 2009/0224502, September 2009.

[33] Zhongjie Li, Lei Zuo, Jian Kuang, and George Luhrs, "Energy-Harvesting Shock Absorber with Mechanical Motion Rectification," *Smart Materials and Structures*, 2012.

[34] Lei Zuo et al., "High-Efficiency Energy Generator for Harnessing Mechanical Vibration Power," 2013/0008157, January 10, 2010.

[35] Ronald Goldner and Peter Zerigian, "Electromagnetic Linear Generator and Shock Absorber," 6,952,060, Oct 4, 2005.

[36] Michael Chen, Christos Papageorgiou, Frank Scheibe, Fu-Cheng Wang, and Malcolm Smith, "The Missing Mechanical Circuit Element," *IEEE Circuits and Systems*, vol. 9, no. 1, 1st quarter 2009.

[37] Malcolm Smith, "Force-Controlling Mechanical Device," 7,316,303, Jan 8, 2008.

[38] Fu-Cheng Wang, Mao-Sheng Hsu, Wei-Jiun Su, and Tz-Chain Lin, "Screw Type Inerter Mechanism," 2009/0108510, Apr 30, 2009.

[39] Fu-Cheng Wang and Tz-Chian Lin, "Hydraulic Inerter Mechanism," 2009/0139225, June 4, 2009.

[40] Matthew Sadiku, *Elements of Electromagnetics*, 3rd ed.: Oxford University Press, 2001.

[41] Robert Crease, *The Great Equations: Breakthroughs in Science from Pythagoras to Heisenberg*, 1st ed.: Norton & Company, 2008.

[42] "Einstein the Greatest," *BBC News*, November 1999.

[43] Patrick Mcfall, "Brainy young James wasn't so daft after all," *The Sunday Post*, April 2006.

[44] Ian Johnston, "Spotlight shines on Maxwell's genius 175 years on through special events," *The Scotsman*, April 2006.

[45] Myron Evans, *Modern Nonlinear Optics, Volume 2*, 2nd ed.: Wiley-Interscience, 2001.

[46] A. E. Fitzgerald, *Electric Machinery*, 3rd ed.: McGraw-Hill Higher Education, 2003.

[47] David Meeker, "Finite Element Mathod Magnetics - User's Manual," 2006.

[48] Alex Goldman, *Handbook of Modern Ferromagnetic Materials*.: Springer, 1999.

[49] E. Wohlfarth and K. Buschow, *Ferromagnetic Materials: a Handbook on the Properties of Magnetically Ordered Substances*.: Elsevier, 1980.

[50] Manitou, "1998 Service Manual".

[51] Manitou, "2012 Service Manual".

[52] HobbyKing. [Online]. http://www.hobbyking.com

[53] Precision Planetary Gears. [Online]. http://www.planetary-precision.com

[54] KSS. [Online]. http://www.kssballscrew.com

[55] The Precision Alliance. [Online]. http://www.tpa-us.com/

[56] Clement Nagode, Mehdi Ahmadian, and Saied Taheri, "Rotational Energy Harvesting Systems for Unpowered Freight Rail Cars," in *Rail Transportation Division Fall Technical Conference*, Roanoke, VA, 2010.

[57] Carl Nelson, Stephen Platt, Sean Hansen, and Mahmood Fateh, "Power Harvesting for Railroad Track Safety Enhancement using Vertical Track Displacement," in *SPIE*, San Diego, CA, 2009.

[58] Abolfazl Pourghodrat, Carl Nelson, Kyle Phillips, and Mahmood Fateh, "Improving an Energy Harvesting Device for Railroad Safety Applications," in *SPIE*, San Diego, CA, 2011.

[59] John Wang, Gopinath Penamalli, and Lei Zuo, "Electromagnetic energy harvesting from train induced railway track vibrations," in *IEEE/ASME International Conference on Mechatronics and Embedded Systems and Applications*, Suzhou, China, 2012.

[60] Timken. [Online]. http://www.timken.com

[61] MatWeb. [Online]. http://www.matweb.com

[62] Stieber Clutch. [Online]. http://www.stieber.de

[63] Redwood Plastics. [Online]. http://www.redwoodplastics.com

[64] Clement Nagode, Mehdi Ahmadian, and Saied Taheri, "Axle Generator for Freight Car Electric Systems," in *Joint Rail Conference*, Philadelphia, PA, 2012.

[65] Innowattech Ltd. [Online]. http://www.innowattech.co.il/

[66] Haim Abramovich, Eugeny Harash, Charles Milgrom, Uri Amit, and Lucy Edery Azulay, "Power Harvesting from Railway: Apparatus, System and Method," 7,812,508, Ocotber 12, 2010.

[67] Carl Nelson, Stephen Platt, Dave Albrecht, Vedvyas Kamarajugadda, and Mahmood Fateh, "Power Harvesting for Railroad Track Health Monitoring Using Piezoelectric and Inductive Devices," in *SPIE*, 2008.

[68] Sean Hansen, Abolfazl Pourghodrat, Carl Nelson, and Mahmood Fateh, "On-track Testing of a Power Harvesting Device for Railroad Track Health Monitoring," in *SPIE*, San Diego, CA, 2010.

[69] Gizmag. [Online]. http://www.gizmag.com

[70] Clement Nagode, Mehdi Ahmadian, and Saied Taheri, "Energy Harvesting Systems to Power Onboard Railroad Equipment," in *Rail Transportation Division Fall Technical Conference*, Minneapolis, Minnesota, 2011.

Appendix A FEMM Lua Script

```
showconsole()
clearconsole()

mydir="./"
input_file = "MT2_Start"
open(mydir .. input_file .. ".FEM")
mi_saveas(mydir .. "temp.FEM")
mi_seteditmode("group")

nfinal=40
picture=0
amp=0.25
freq=1
nosc=1
pos=0

mi_selectgroup(1)
mi_movetranslate(0,pos)

for n=0,nfinal do
      mi_analyze()
      mi_loadsolution()

      mo_addcontour(0,2.875)
      mo_addcontour(0.52,2.875)
      BnTotal = mo_lineintegral(0)
      mo_clearcontour()

      mo_addcontour(1,2.875)
      mo_addcontour(1.125,2.875)
      BnTotal2 = mo_lineintegral(0)
      mo_clearcontour()

      print(n/(nfinal/nosc),pos,BnTotal,BnTotal2)

      if (picture==1) then
            mo_showdensityplot(1,0,2,0.001,"mag")
            mo_savebitmap(mydir .. input_file .. "_" .. n ..".bmp")
      end

      if (n<nfinal) then
            d      =      amp*cos(360*freq*(n+1)/(freq*nfinal/nosc))-
amp*cos(360*freq*n/(freq*nfinal/nosc))
            mi_selectgroup(1)
            mi_movetranslate(0,d)
      end

      pos=pos+d
end

mo_close()
mi_close()
```

This page intentionally left blank.

Appendix B Average Values: RMS vs. Mean

In many cases, the result of a test is summed up into one number such as, for example, the average power over a cycle. For the average power to be truly meaningful, it needs to be equal to the power that would yield the same amount of energy:

$$P_{avg} \cdot \Delta t = E = \int_{t_i}^{t_f} P(t)\,dt$$

This clearly shows that to represent the power generated over a period of time, the arithmetic mean needs to be used:

$$P_{avg} = \frac{1}{t_f - t_i}\int_{t_i}^{t_f} P(t)\,dt \qquad \text{or} \qquad P_{avg} = \frac{1}{N}\sum_{i=1}^{N} P_i \text{ , for a discrete set of data}$$

This is different from the RMS value of the power, which is expressed as:

$$P_{RMS} = \sqrt{\frac{1}{t_f - t_i}\int_{t_i}^{t_f} \left[P(t)\right]^2 dt} = \sqrt{\frac{1}{N}\sum_{i=1}^{N} P_i^2}$$

The RMS value of the power, which would be expressed in Watt RMS, has no physical significance and should not be used. Often, especially in audio equipment, the output power is confusingly expressed in W_{RMS}. It does not mean the RMS value of the power, but refers to the power calculated from the RMS value of the voltage with a continuous sine wave signal, as prescribed by the Federal Trade Commission (FTC).

It is important to note that the RMS value of any signal is always greater than or equal to the mean value, both being equal only for a constant signal. So, in almost every case, using the RMS value will overestimate the actual power generated. On the other hand, the RMS values should be used for voltages and currents (in purely resistive loads) as they represent the constant signal that yields the same time-averaged power dissipation, i.e. $P_{avg} = V_{RMS} \cdot I_{RMS} = \frac{{V_{RMS}}^2}{R} = R \cdot {I_{RMS}}^2$

This can be easily explained by going back the definition of the average power:

$$P_{avg} = \frac{1}{t_f - t_i}\int_{t_i}^{t_f} P(t)\,dt = \frac{1}{t_f - t_i}\int_{t_i}^{t_f} V(t)\cdot I(t)\,dt = \frac{1}{t_f - t_i}\int_{t_i}^{t_f} V(t)\cdot\frac{V(t)}{R}\,dt = \frac{1}{R}\frac{1}{t_f - t_i}\int_{t_i}^{t_f}\left[V(t)\right]^2 dt$$

$$\Rightarrow P_{avg} = \frac{{V_{RMS}}^2}{R} \text{ since, by definition, } V_{RMS} = \sqrt{\frac{1}{t_f - t_i}\int_{t_i}^{t_f}\left[V(t)\right]^2 dt}$$

This page intentionally left blank.

Appendix C Harvester Mounting Guidelines

Mounting guidelines

The energy harvesting system is meant to convert part of the energy normally dissipated in the damping system (friction wedges) into useful electrical energy. The system is designed to be placed inside a D-5 spring. However, for initial field testing, it is more convenient to install the system on the side of the bogie (Figure 1). The red parts are examples of the brackets that need to be added to the side frame and the bolster to mount the system.

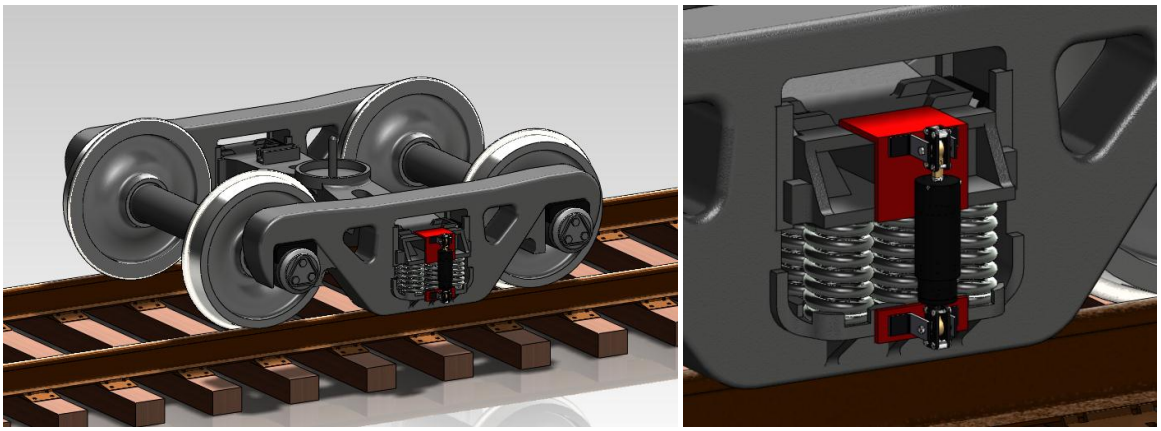

Figure 1: Overview

The system needs to be mounted parallel to the suspension. $d_{harvester}$, the distance between the mounting points (i.e. between the centers of the ball joints) is defined by the following equation:

$$d_{harvester} = l_{spring} + 6in \pm 0.25in$$

with l_{spring} the (loaded) spring height (Figure 2).

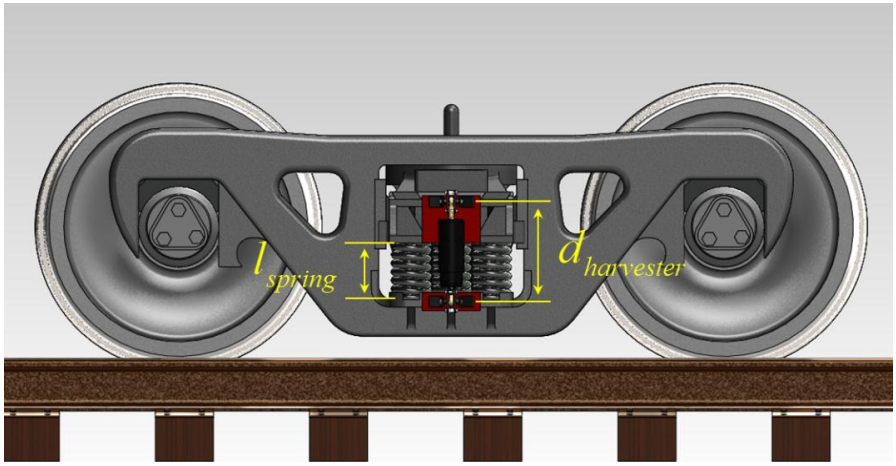

Figure 2: Mounting dimensions